Simon Reiter

Implementation, Operation and Monitoring of the Data Acquisition System of the Belle II Pixel Detector during Physics Data Taking in 2019 – 2022

Dissertation



Implementation, Operation and Monitoring of the Data Acquisition System of the Belle II Pixel Detector during Physics Data Taking in 2019 - 2022





Implementation, Operation and Monitoring of the Data Acquisition System of the Belle II Pixel Detector during Physics Data Taking in 2019 – 2022

Implementierung, Betrieb und Überwachung des Auslesesystems für einen Silizium-Pixeldetektor am Belle II Experiment während der Datennahme von 2019-2022

Simon Patrik Reiter

INAUGURALDISSERTATION zur Erlangung des Doktorgrades

am

II. Physikalischen Institut

des

Fachbereich 07 — Mathematik und Informatik, Physik, Geographie *der* Justus-Liebig-Universität Gießen

(2023)

Dekan: Prof. Dr. Stefan Björn Hennemann Gutachter: apl. Prof. Dr. Jens Sören Lange Gutachter: Prof. Dr. Claudia Höhne

This document was typeset with Lua LTEX 1.17 using the memoir class. The text is set in Libertine, Biolinum, and LibertineMono.

Selbstständigkeitserklärung

Hiermit versichere ich, die vorgelegte Thesis selbstständig und ohne unerlaubte fremde Hilfe und nur mit den Hilfen angefertigt zu haben, die ich in der Thesis angegeben habe. Alle Textstellen, die wörtlich oder sinngemäß aus veröffentlichten Schriften entnommen sind, und alle Angaben die auf mündlichen Auskünften beruhen, sind als solche kenntlich gemacht. Bei den von mir durchgeführten und in der Thesis erwähnten Untersuchungen habe ich die Grundsätze gute wissenschaftlicher Praxis, wie sie in der "Satzung der Justus-Liebig-Universität zur Sicherung guter wissenschaftlicher Praxis" niedergelegt sind, eingehalten. Entsprechend § 22 Abs. 2 der Allgemeinen Bestimmungen für modularisierte Studiengänge dulde ich eine Überprüfung der Thesis mittels Anti-Plagiatssoftware.

Datum

Unterschrift

Kurzdarstellung

Die vorliegende Arbeit fasst meinen Beitrag zur Echtzeit-Datennahme des Pixeldetektors des Belle II Experiments am High Energy Accelerator Research Organization (KEK) in Tsukuba, Japan, zusammen. Das Belle II Experiment befindet sich am SuperKEKB Teilchenbeschleuniger, bei dem Elektronen und Positronen mit einer Schwerpunktsenergie von 7.4 GeV aufeinander treffen. Dabei entstehen kurzlebige Teilchen, die nach nur wenigen Mikrometern zerfallen. Aufgrund einer Lebensdauer von teilweise weniger als einer Picosekunde, ist die örtliche Auflösung für die vollständige Rekonstruktion des Zerfalls von entscheidender Rolle. Der Pixeldetektor (PXD), speziell für dieses Experiment entwickelt, besteht aus zwei Lagen, die in einem Abstand von 14 mm und 22 mm zylindrisch um den Wechselwirkungspunkt angebracht sind. Aufgrund der geringen Distanz und der hohen Luminosität von 6×10^{35} cm⁻² s⁻¹, werden viele Untergrundsignale erwartet, die nicht für weitere Auswertungen geeignet sind. Durch den langsamen Ausleseprozess von 20 us werden Signale in bis zu 3 % aller Pixel erwartet. Dies entspricht einer Datenrate von über 17 GB/s. Das einheitliche Datennahmesystem aller anderen Detektoren ist dieser Datenmenge nicht gewachsen, so dass an der Justus-Liebig-Universität in Gießen das ONSEN-System entwickelt wurde. Speziell abgestimmte Hardwarekomponenten wurden ausgewählt und in Kooperation mit dem Forschungszentrum IHEP in Peking eine FPGA-basierte Hardware-Plattform entwickelt. Die Firmware des FPGAs ist so konzipiert, dass die Daten mit Hilfe von Bereichen von Interesse reduziert werden, welche durch eine Echtzeit-Ereignisrekonstruktion ermittelt werden. Bevor die Daten gespeichert werden kann so eine Datenreduktion um den Faktor 30 erfolgen. Während der Datennahme wird der PXD sowie alle Komponenten, die für seinen reibungslosen Betrieb notwendig sind, permanent überwacht. Hier kommt eine Software zum Einsatz, die speziell für Physik- und Beschleuniger-Experimente entwickelt wurde. Das Experimental Physics and Industrial Control System (EPICS) zeichnet sich dadurch aus, dass es eine Vielzahl an Geräten unterstützt und das Überwachen und Steuern in einer graphischen Oberfläche dargestellt werden kann.

Dieses Dokument fasst die erfolgreiche Erweiterung des ONSEN-Systems sowie die Integration in das Belle II Experiment zusammen. Die Firmware wurde ergänzt um Überwachungsmechanismen der Datenverarbeitung, um eine reibungslose Datennahme zu gewährleisten und ihre Effizienz zu maximieren. Die kontinuierliche Berechnung der signalgebenden Pixel sowieso die Erkennung von unvollständigen Datensätzen stellten sich als eine der wichtigsten Entscheidungskriterien raus, die zum Bestimmung der Detektorleistung genutzt werden. Der Integrationsprozess umfasste die nahtlose Einbindung der ONSEN-Firmware in EPICS zum Bereitstellen von Echtzeitinformationen. Dies hat sich für die Fehlersuche und den Betrieb als unschätzbar erwiesen, da sie eine rasche Identifizierung und Lösung potenzieller Probleme ermöglicht. Das ONSEN-System hat erfolgreich an der Erfassung von physikalischen Daten teilgenommen und seine Zuverlässigkeit und Robustheit unter Beweis gestellt. Im Laufe von drei Jahren, zwischen 2019 und 2022, hat ONSEN fast 300 TB an Rohdaten zuverlässig analysiert und verarbeitet. In gesonderten Studien wurden Testdaten mit Datenraten von bis zu 4.8 GB/s generiert und unterbrechungsfrei vom ONSEN-System verarbeitet.

Um die Zuverlässigkeit des gesamten PXD zu erhöhen und den aufwendigen Betrieb des Beschleunigers und des Belle II Experiments so effizient wie möglich zu gestalten, wurde die Integration in EPICS des gesamten PXD überprüft. Neben Fehlerbehebung wurden Mechanismen implementiert, um neuen Anforderungen gerecht zu werden, die erst nach Inbetriebnahme ersichtlich wurden. Durch Auswertung von wiederkehrenden manuellen Arbeitsabläufen wurden diese in automatisierte Prozesse überführt, um die Wiederherstellung nach Modulausfällen und die Kalibrierung der Sensoren zu beschleunigen.

Abstract

This paper summarizes my contribution to the real-time data taking of the pixel detector of the Belle II experiment at the High Energy Accelerator Research Organization (KEK) in Tsukuba, Japan. The Belle II experiment is located at the SuperKEKB particle accelerator, where electrons and positrons collide with a center-of-mass energy of 7.4 GeV. In this process, short-lived particles decay after a few micrometers. Due to a lifetime of sometimes less than a picosecond, local resolution is crucial for the complete reconstruction of the decay. The pixel detector (PXD), specifically designed for this experiment, consists of two layers, spaced 14 mm and 22 mm apart cylindrically around the interaction point. Due to the small distance and the high luminosity of 6×10^{35} cm⁻² s⁻¹, many background signals are expected, which are not suitable for further analysis. Due to the slow readout process of 20 µs, signals are expected in up to 3% of all pixels. This corresponds to a data rate of over 17 GB/s. The uniform data acquisition system of all other detectors cannot cope with this amount of data, so the ONSEN system was designed at the Justus Liebig University in Giessen. Specially matched hardware components were selected and an FPGA-based hardware platform was developed in cooperation with the IHEP research center in Beijing. Fine tuned hardware components were selected and an FPGA-based hardware platform was developed in cooperation with the IHEP research center in Beijing. The FPGA firmware is designed to reduce data using regions of interest determined by real-time event reconstruction. Thus, before data is stored, a data reduction by a factor of 30 can be performed. During data acquisition, the PXD and all components necessary for its smooth operation are constantly monitored. Software specially developed for physics and accelerator experiments is used here. The Experimental Physics and Industrial Control System (EPICS) supports a large number of devices and allows monitoring and control information to be displayed in a graphical user interface.

This document summarizes the successful extension of the ONSEN system and its integration into the Belle II experiment. Monitoring mechanisms for data processing were added to the firmware to ensure smooth data acquisition and maximize efficiency. The continuous calculation of signaling pixels and detection of incomplete data sets turned out to be one of the most important criteria used to determine the detector performance. The integration process included seamless integration of the ONSEN firmware into EPICS to provide real-time information. This has proven invaluable for troubleshooting and operations by enabling rapid identification and resolution

of potential problems. The ONSEN system has successfully acquired physics data, demonstrating its reliability and robustness. Over three years, between 2019 and 2022, ONSEN reliably analyzed and processed nearly 300 TB of raw data. In separate studies, test data with data rates of up to 4.8 GB/s were generated and processed by the ONSEN system without interruption.

To increase the reliability of the entire PXD and to make the operation of the accelerator and the Belle II experiment as efficient as possible, the integration into EPICS of the entire PXD was reviewed. In addition to troubleshooting, mechanisms were implemented to address constraints that became apparent only after commissioning. By evaluating recurring manual workflows, these were transformed into automated processes to speed up recovery from module failures and sensor calibration.

Contents

1	Introduction						
2	Theoretical Considerations						
	2.1	The Standard Model of Particle Physics	3				
	2.2	The Standard Model Lagrangian	5				
	2.3	CP Violation	11				
3	The Belle II Experiment						
	3.1	The SuperKEKB Accelerator	19				
	3.2	The Belle II Detector	21				
		3.2.1 KLM – The K_L^0 and Muon Detector	22				
		3.2.2 ECL – The Electromagnetic Calorimeter	23				
		3.2.3 ARICH – The Aerogel Ring-Imaging Cherenkov Detector .	24				
		3.2.4 TOP – The Time of Propagation Detector	24				
		3.2.5 CDC – The Central Drift Chamber	25				
		3.2.6 VXD – The Vertex Detector	25				
	3.3	PXD – The Pixel Detector	26				
4	Data Acquisition 33						
	4.1	Belle II DAO Overview	33				
		4.1.1 Trigger System	33				
		4.1.2 Data Flow	34				
	4.2	PXD Data Handling	36				
	43	Online Selection Node	38				
	1.5	4 3 1 Hardware	40				
		432 Firmware	41				
		4 3 3 Software	46				
		434 Status	47				
		4.3.5 Future Development	48				
5	Slow Control 49						
5	5 1	The EPICS Ecosystem	49				
	5.2	ONSEN Slow Control	51				
	5.2	ONSEN Run Control	54				

References							
7	Summary and Outlook						
	6.2	Physic	s Results	84			
	6.1	Instrur	nental Results	76			
6	Operation Results						
		5.5.5	Calibration IOC	72			
		5.5.4	Data Transfer IOC	71			
		5.5.3	Local Data Quality Monitor	70			
		5.5.2	Local DAQ Control IOC	69			
	0.0	5.5.1	PXD Shifter IOC	68			
	5.5	PXD C	alibration Framework	66			
		5 4 10	PXD Run Control Undate	65			
		5.4.0	DHH Sequence Undate	62			
		5.4.7	PS Control Update	60			
		5.4.6	Clear Channel Handling	60			
		5.4.5	DHP Register Monitor	59			
		5.4.4	Test Pattern Upload	59			
		5.4.3	Module Recovery	58			
		5.4.2	DHP Link Recovery	57			
		5.4.1	Overview	55			
	5.4 PXD Slow Control						

Chapter 1

Introduction

Particle physics experiments are collaborative efforts by researchers investigating nature at small scales. The conduct of an experiment is motivated by a preceding, promising physics theory or exploration of the limits of an established model. The planning and construction of these machines is extremely complex, and experimental physicists design new detectors and, in many cases, dedicated particle accelerators suitable for the task. In order to begin taking data and obtaining long-awaited results, all components of the experiment must be assembled and working, including detectors and data acquisition systems.

There are two types of high energy physics experiments: The ATLAS and CMS experiments at the LHC proton-proton collider utilize large instruments to produce extremely high energy particle beams. Using this method, they can investigate physics processes not possible at lower energies and discover new, very massive particles. In contrast, intensity frontier experiments use very precise detectors and intense particle beams to investigate rare processes. This category includes the Belle experiment at the KEKB electron-positron collider in Tsukuba, Japan.

The Belle experiment concluded in 2010 and plans for an upgrade were approved. This successor Belle II is now the only second generation B-factory. With the upgrade of the experiment comes an upgrade of the KEKB accelerator. The luminosity will be increased by a factor of 30 and a newly developed vertex detector based on a state-of-the-art technology will increase the vertex resolution. This upgrade will amplify the background rate and additional data must be stored.

When the collider reaches its design luminosity, the pixel detector (PXD) will be dominated by background hits and produce data rates of up to 20 GB/s due to the short distance to the interaction point. It was necessary to create a dedicated data acquisition (DAQ) system for the PXD, running concurrently with the Belle II DAQ system, to handle and reduce these high data rates. Therefore the Belle II group at the University of Gießen developed the Online Selection Node (ONSEN) system, using a hardware platform designed at the IHEP in Beijing, China.

The Belle II experiment started operation in 2019, and the ONSEN system proved its capabilities for the first time. This thesis describes the improvement of the firmware and software aspects of the ONSEN system to increase data taking efficiency. Additionally the overall performance of the ONSEN system during the first three years is presented.

For the operation and optimization of the PXD and its accompanying components a sophisticated control and monitoring system is necessary. As part of the thesis to maximize data taking efficiency, system behavior was analyzed in various situations and the software was revised and extended. Individual workflows and handling unpredictable system failures were automated to reduce the downtime.

Chapter 2

Theoretical Considerations

An overview of modern particle physics is provided in this chapter, along with a discussion of the physics of interest to the Belle II experiment. This section describes the forces and particles that make up the *Standard Model* (SM), beginning with an introduction to fundamental particles and their categorization before discussing the most relevant terms of the SM Lagrangian and its implications for *charge conjugation and parity violation*. A brief description of Belle II's method of measuring CP-violations is provided in the second section.

2.1 The Standard Model of Particle Physics

According to current understanding, particle physics examines *elementary particles*, which are nearly point-like objects that cannot be subdivided into smaller components. As a result of Yang and Mills [1], Glashow [2] and Weinberg and Salam's work [3] beginning in the middle of the 20th century, a successful theoretical framework, the Standard Model (SM), has been developed to describe these particles and their interactions in the last decades. With each experimental validation, the SM developed further in the following decades, such as in 1973 when the weak neutral current was discovered [4]. The SM is based on quantum field theories (QFT), theoretical frameworks that combine classical field theories, quantum mechanics, and special relativity. A number of those have been developed on the basis of the two most fundamental QFTs over the past years, quantum chromodynamics (QCD) and quantum field theories with an underlying $SU(3)_c \times SU(2)_L \times U(1)_Y$ gauge group¹.

This unified theoretical building is essential to describe three out of the four forces in nature:

¹Here c is the color charge, L the chiral symmetry and Y the hyper-charge. First gauge theory proposed by W. Pauli in 1941 with an invariant Lagrangian under local transformations [5].

Strong Force

Described by QCD, it contains quarks into proton, neutron, and other hadron particles. It binds neutrons and protons together in atomic nuclei, where it is referred to as the nuclear force.

Weak Force

Described by quantum electrodynamics, it is the mechanism of interaction between subatomic particles that are responsible for the radioactive decay of atoms.

Electromagnetism

Described by quantum electrodynamics, it governs the interaction between atoms and molecules and occurs when particles with electric charges are in contact with electromagnetic fields.

The SM does not describe the most recognizable force, gravitation. As a result of construction, the existence of gravitational force-carrying particles in the Standard Model is valid, but describing phenomena in the Standard Model cannot be accomplished. Due to the small scales and masses resulting from distance and time evolution, gravity is suppressed by many orders of magnitude and may be ignored. It is shown in figure 2.1 how elementary particles and forces are connected in the Standard Model.

In figure 2.1, particles are classified into two kinds. Fermions are particles with spin 1/2 and are responsible for the matter around us. They can be divided into quarks (blue), which carry a color charge and are sensitive to the strong force, and leptons (red). The bosons, on the other hand, have integer spin. Most are resulting from excitations in the associated gauge fields and therefore called gauge bosons (yellow). The strong force is mediated by the gluon (g), the electromagnetic force is mediated by the photon (γ) , and the weak force is carried by the W^{\pm} and Z^{0} boson. The SM does not use the Higgs boson (\mathcal{H}) to mediate any force, but instead to acquire mass. The Higgs field is responsible for introducing spontaneous symmetry breaking, which is a phenomenon introduced by a complex scalar field that has a non-zero vacuum expectation value in the SM².

²It is the average value that an operator acquires in a vacuum, for example, the state with the lowest energy in the universe.



Figure 2.1: A schematic overview of the fundamental constituents of the Standard Model of Particle Physics.

2.2 The Standard Model Lagrangian

In this section, we will briefly discuss each term of the SM Lagrangian formalism, which is used to represent the mathematical description of the theories unified within the SM.

$$\mathcal{L}_{SM} = \mathcal{L}_{kin} + \mathcal{L}_{EW} + \mathcal{L}_{QCD} + \mathcal{L}_{H} + \mathcal{L}_{YU}$$
(2.1)

In the first term, $\mathcal{L}_{kin},$ the self-interactions of the gauge bosons are described. It can be written as

$$\mathcal{L}_{kin} = -\frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{2} tr \left(W_{\mu\nu} W_{\mu\nu} \right) - \frac{1}{2} tr \left(G_{\mu\nu} G^{\mu\nu} \right), \tag{2.2}$$

where traces (tr) run over the gauge field tensor of SU(2) and the gluon field tensor of SU(3) indices of W and G respectively. $B_{\mu\nu}$ represents the gauge field tensor of the weak hyper charge.

Electroweak unification of the SM is described by the second term, \mathcal{L}_{EW} , which correlates to the symmetry group $U(1) \times SU(2)_L^{-3}$.

$$\mathcal{L}_{EW} = \sum_{i} \overline{\psi}_{i} \gamma^{\mu} \left(i \partial_{\mu} - \frac{1}{2} g' Y_{W} B_{\mu} - \frac{1}{2} g \tau W_{\mu}^{a} \right) \psi_{i}$$
(2.3)

A left-handed fermion doublet is depicted as a Dirac spinor in the above equation, which is connected to a vector representation of the Lorentz group by the Dirac matrices γ_{μ} . As a consequence, the derivatives ∂_{μ} introduce a kinetic term for fermions. g' and g are, respectively, coupling constants for electromagnetic and weak fields. The weak hyper charge Y_W is determined as the generator of the U(1) group acting on the U(1) gauge field B_{μ} . It can be related to the electric charge Q and the third component of the weak isospin I_3 by

$$Q = I_3 + \frac{1}{2}Y_W$$

Pauli matrices τ act as infinitesimal generators of the SU(2) group and contain the isospin charges of particles interacted with the weak force as their eigenvalues. *Electroweak unification* is achieved by using gauge bosons, in which a weak isospin triplet (W^a_μ with a = 1, 2, 3) and a singlet (B_μ) act as mediators between electromagnetic and weak forces [6, 3].

This results in the following couplings for each boson:

photon (γ)

Are involved in the electromagnetic interaction, since they couple to the electric charge. In the absence of mass and charge, the photon itself becomes the basis for electromagnetic interactions with an indefinite range of application. Neutral particles are not affected.

charged W^{\pm} boson

Mediates the charged current of the weak force. Following the SU(2) group, the W^{\pm} is maximally charged and parity-violating. Therefore, it couples only with left-handed particles and right-handed antiparticles. At present, only the exchange of the W^{\pm} bosons is capable of a flavor change from up-type quarks to down-type quarks and leptons to its neutrino pair in the SM. Such flavor changes in the lepton sector are limited to individual families. Flavor changes in the quark sector have no such boundaries.

neutral Z^0 boson

Mediates the neutral current of the weak force. The interaction is similar to the electromagnetic interaction. Since the W^{\pm} and the Z^0 have very large masses, the weak interaction is short-ranged and has a similar coupling strength

 $^{^{3}}L$ indicates a coupling only to left-handed fermions, where the spin points in the opposite direction of motion.

to the electromagnetic interaction. For this reason, a flavor change via the neutral current is not forbidden in the SM, but strongly suppressed via the GIM mechanism [6].

The term \mathcal{L}_{QCD} is introduced in order to describe the strong interactions in the SM.

$$\mathcal{L}_{QCD} = \sum_{j} \overline{\psi}_{j} \gamma^{\mu} \left(i g_{s} G_{A,\mu} T_{A,ij} - m \delta_{ij} \right) \psi_{j}$$
(2.4)

 g_s refers to the strong force coupling constant, $G_{A,\mu}$ refers to the gluon field, T_a refers to the symmetry generator, and m refers to the mass. There are two indices i and j, which indicate quarks' and gluon's color and anti-color charge. The color charge or only referred to as color is a quantum number, which can be either red, green or blue. Its inverse (anti-color) is also possible for each value. Due to the non-abelian nature of the underlying symmetry group SU(3), gluons always carry color and anti-color at the same time, which allows them to interact with one another. When it comes to quarks, they only carry either color, or anti-color, depending on their type.

The strong force binds elementary particles to hadrons. Depending on their spin they can be classified as mesons with integer spin due to its composition of quark and anti-quark or baryons with half spin⁴ composed of 3 (anti-) quarks. The combination of the color of all components is always zero. This behavior is called *confinement* [7, 8] and results in a prohibited observation of colorful particles including unbound quarks. The SM allows bound states with a higher number of constituents. The first observation of such is the X(3872), seen by the Belle experiment in 2003 [9] and later validated by the BaBar and BESIII experiment [10, 11].

As described in the Higgs-term of the Lagrangian, the Higgs-boson contributes kinetic energy to the system along with its interactions with the other force carriers. It can be described like this:

$$\mathcal{L}_{H} = \left[\left(\partial_{\mu} - igW_{\mu}^{a}t^{a} - ig'Y_{\phi}B_{\mu} \right)\phi \right]^{2} + \mu^{2}\phi^{\dagger}\phi - \lambda(\phi^{\dagger}\phi)^{2}$$
(2.5)

The kinetic term is given by ∂_{μ} . The coupling to the weak and electromagnetic force is contributed via g and g'. The Higgs field is the complex scalar ϕ , which is described in the following with superscripts representing the electric charge.

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}$$

The parameter λ and μ^2 have to be positive to allow spontaneous symmetry breaking.

⁴The half spin values follow the formula (2n + 1)/2. No integer spin values are possible.

The last term of the Lagrangian is the Yukawa term. It is responsible for the mass acquisition of quarks. In the equation below $G_{u,d}$ are 3×3 matrices with complex elements each representing the quark generations, where u and d denote up- or down-type quarks. The scalar fields ϕ couple left to right handed quark fields $D_{L,B}$.

$$\mathcal{L}_{YU} = G_u^{ij} \overline{U}_L \phi^0 U_R - G_u^{ij} \overline{D}_L \phi^- U_R + G_d^{ij} \overline{U}_L \phi^+ D_R + G_d^{ij} \overline{D}_L \phi^0 D_R \qquad (2.6)$$

Due to the non-zero vacuum expectations on the scalar Higgs fields ϕ , the mass matrices (MU_{ij}) and (MD_{ij}) for up-type and down-type quarks can be expressed as follows:

$$MU_{ij} = \frac{1}{\sqrt{2}} G_u^{ij} v^0$$
 and $MD_{ij} = \frac{1}{\sqrt{2}} G_d^{ij} v^0$ (2.7)

Through diagonalization with the unitarity matrices $TU_{L,R}$ and $TD_{L,R}$, the mass matrices get non-zero entries on their main diagonals. In addition, these matrices transform the weak eigenstate bases of left-handed and right-handed quark fields into mass eigenstate bases. Rather than expressing quark interaction with gauge fields as a weak interaction, the neutral current term remains invariant, while the charge current term changes. GIM is based on this concept, which is why flavor-changing neutral currents are not present at the tree level in SM. The charged currents looks like this:

$$I_{\nu}^{+} = \overline{U}_{L} \gamma^{\nu} D_{L} = \overline{U}_{L}^{M} \gamma^{\nu} T U_{L} T D_{L}^{\dagger} D_{L}^{M}$$

$$I_{\nu}^{-} = \overline{D}_{L} \gamma^{\nu} U_{L} = \overline{D}_{L}^{M} \gamma^{\nu} T D_{L} T U_{L}^{\dagger} U_{L}^{M}$$
(2.8)

As a result of expressing the charged currents in the mass eigenstates (^M), one can express the transformation matrix $TU_L TD_L^{\dagger}$ as the Cabibbo-Kobayashi-Maskawa (CKM) unitary matrix V_{CKM} .

. .

$$V_{CKM} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}$$
(2.9)

As one observes from comparing the charged current interaction with its charge and parity (CP) conjugated counterpart, the interactions are only equal when

$$V_{CKM_{ij}}^* = V_{CKM_{ij}}$$

is satisfied. The conjugation is made by the exchange of all internal quantum numbers (C) and inversion of the handedness (P). According to Kobayashi and Maskawa, the CKM quark-mixing matrix contains a complex phase when at least three generations



Figure 2.2: Matrix elements of the CKM matrix are visualized as intensity lines between quark families. The thickness correlates to the coupling strength. CKM prefers transitions between quarks belonging to the same family, like u to d and vice versa.

of quarks are present [12]. The result of this phase is that the matrix elements i and j of V_{CKM} differ from those of V_{CKM}^* , which constitutes the only source of violation of CP in the SM.

In the CKM matrix, the corresponding matrix element couples the up- and downquarks through charged current interactions mediated by the W^{\pm} bosons. Feynman diagrams contain matrix elements that describe the strength of the coupling between the charged weak boson and the quark flavors. Figure 2.2 illustrates the strength of this coupling. The coupling strength is indicated by the line thickness between different quark flavors. A higher degree of coupling results in a higher likelihood of a transition occurring. The unitarity conditions the CKM matrix can be expressed like this:

$$\sum_{i} (V_{CKM})_{ij} (V_{CKM})_{ik}^* = \delta_{jk}$$
(2.10)

$$\sum_{j} (V_{CKM})_{ij} (V_{CKM})_{kj}^* = \delta_{ik}$$
(2.11)

These conditions can be combined resulting in six equations that allow a different perception.

$$\begin{split} V_{ud}^*V_{us} + V_{cd}^*V_{cs} + V_{td}^*V_{ts} &= 0\\ V_{ud}V_{cd}^* + V_{us}V_{cs}^* + V_{ub}V_{cb}^* &= 0\\ V_{us}^*V_{ub} + V_{cs}^*V_{cb} + V_{ts}^*V_{tb} &= 0\\ V_{td}V_{cd}^* + V_{ts}V_{cs}^* + V_{tb}V_{cb}^* &= 0\\ V_{td}V_{ud}^* + V_{ts}V_{us}^* + V_{tb}V_{ub}^* &= 0\\ V_{ud}V_{ub}^* + V_{cd}V_{cb}^* + V_{td}V_{tb}^* &= 0 \end{split}$$

A triangle in the complex plane can be interpreted as a result of all of these equations. It is possible to calculate CP violation in the CKM matrix by measuring the angles of these triangles. In figure 2.2, the coupling strength between quark flavors is represented by lines with different thicknesses. The numerical difference between thick and thin lines is a few orders of magnitude. As a result of the first four equations above, the triangle will appear distorted, with very small angles. Their corresponding decay channels are difficult to measure. In contrast, the last two equations produce triangles with larger angles due to the fact that the matrix elements all have the same magnitude. In *B*-factories like Belle II, where the produced *B* mesons contain a bottom quark, the triangle represented by the last equation contains bottom quark transitions in each terms and is of higher interest. The triangle shown in figure 2.3 highlights some of the measurable decay channels. The angles of the triangle are named alpha, beta, and gamma following the Belle naming scheme and correspond to the matrix elements of the CKM matrix.

$$\begin{aligned} \alpha &= \arg\left(-\frac{V_{td}V_{tb}^*}{V_{ud}V_{ub}^*}\right), \quad \beta = \arg\left(-\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*}\right), \\ \gamma &= \arg\left(-\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*}\right) \end{aligned}$$
(2.12)

Observable magnitudes can be used as a gauge of the extent of CP violation: No CP violation would result in $\alpha = \gamma = 1$ and $\beta = 0$. On the other hand, CP violation can lead to significant differences between observables. Possible main decay modes for each angle, which can be measured at *B*-factories like Belle II, are

$$\begin{split} &\alpha \!\!: \; B^0 \to \pi^\pm \pi^\mp, \; \rho \pi^0, \ldots \\ &\beta \!\!: \; B^0 \to J/\Psi K^0_S, \; D^* \overline{D}^*, \ldots \\ &\gamma \!\!: \; B^0 \to D^{*\pm} \pi^\mp, \; D^\pm K^\mp, \ldots \end{split}$$



Figure 2.3: The unitarity triangle in the complex plain with its dependency to the CKM matrix elements.

2.3 CP Violation

According to the SM, CP-violations can be classified into two types – direct and indirect. Direct CP violation results in different branching fractions for a particle and its CP conjugated counterpart due to different decay amplitudes. The only source of CP-violation in charged meson decay is this method. A direct CP-asymmetry A_{CP}^{dir} for charged B mesons can be expressed as the difference in decay rates between B^+ and B^- normalized to their total decay rate.

$$A_{CP}^{dir} = \frac{Br(B^- \to f^-) - Br(B^+ \to f^+)}{Br(B^- \to f^-) + Br(B^+ \to f^+)}$$
(2.13)

An example for direct CP-violation is shown in the Feynman diagram in figure 2.4.



Figure 2.4: Feynman diagram of the decay $B^0_d \to K^+\pi^-.$

Indirect CP violation occurs as a result of meson mixing. Here a meson containing quark and anti-quark is generated with definite quark flavors, but propagates with distinct mass. It is possible to observe a flavor eigenstate, which is different from the initial flavor eigenstate, after a certain amount of time has elapsed. As a result of this change in flavor eigenstate, the particle oscillates into its own anti-particle and back. As shown in figure 2.5, such oscillation occurs in the mixed system of B^0 and B_s^0 . The mixing frequency is significantly higher in the B_s^0 system due to the stronger CKM-coupling between the *s* and *b* quark. As a result of the flavor change having to skip two generations, the matrix element for B^0 mixing is much smaller.



Figure 2.5: Oscillation probabilities for the B^0 (a) and B_s^0 (b) system in dependence of the decay time t.

The eigenstates of B^0 mixing are differentiated for light (*L*) and heavy (*H*) due to slightly different masses. They are given by

$$B_L = p|B^0\rangle + q|\overline{B}^0\rangle \tag{2.14}$$

$$B_H = p|B^0\rangle - q|\overline{B}^0\rangle \tag{2.15}$$

With the application of the charge operator particles are converted to antiparticles and vice versa. The parity operator introduces a negative sign. The result would look like this:

$$CP(p|B^{0}\rangle + q|\overline{B}^{0}\rangle) = -(p|B^{0}\rangle + q|\overline{B}^{0}\rangle)$$
(2.16)

0

$$CP(p|B^{0}\rangle - q|\overline{B}^{0}\rangle) = (p|B^{0}\rangle - q|\overline{B}^{0}\rangle)$$
(2.17)

These equations can be interpreted that the mass eigenstates are identical to the CP eigenstates. This results in the propagation of a definite CP eigenstate in that state and could not be observed with another value at a later time. To measure a changed value the absolute values of p and q must not be equal. While this violation is only allowed during mixing, it is called *indirect CP violation*.

Another form of direct CP violation can be observed when a particle and its CP conjugate decay to a final state f and the CP conjugated final state \overline{f} . The magnitude of CP violation is given by

$$\lambda = \frac{q}{p} \frac{A_f}{\overline{A}_f} \tag{2.18}$$

In case of B mesons, while one B can be determined unambiguously, the other B decays into a CP eigenstate. In such case the decay rates can be expressed as

$$f_{\pm}(\Delta t) = \frac{e^{-|\Delta t|/\tau_{B^0}}}{4\tau_{B^0}} \left[1 \pm \frac{2\operatorname{Im}\lambda}{1+|\lambda|^2}\sin(\Delta m_d\Delta t) \mp \frac{1-|\lambda|^2}{1+|\lambda|^2}\cos(\Delta m_d\Delta t) \right]$$
(2.19)

Here τ_{B^0} is the B^0 lifetime and Δm_d the mass difference between the two B mass eigenstates [13]. The sign for sine and cosine terms is chosen for CP odd final states, for even ones they must be inverted. In order to visualize the possible CP violation, the asymmetry between $f_{\perp}(\Delta t)$ and $f_{-}(\Delta t)$ can be constructed.

$$\mathcal{A}(\Delta t) = \frac{f_+(\Delta t) - f_-(\Delta t)}{f_+(\Delta t) + f_-(\Delta t)},$$
(2.20)

which simplifies to

$$\mathcal{A}(\Delta t) = S \sin(\Delta m_d \Delta t) - C \cos(\Delta m_d \Delta t) \tag{2.21}$$

with the amplitudes

$$S = \frac{2 \operatorname{Im} \lambda}{1 + |\lambda|^2} \qquad \qquad C = \frac{1 - |\lambda|^2}{1 + |\lambda|^2}$$

which are related via the following, where θ is the phase of λ .

$$\left(\frac{S}{\sin\theta}\right)^2 + (C)^2 = 1$$

After fitting of S and C their values can be related to physical quantities. In case the contributing amplitudes carry the same weak phase the absolute amplitudes (\mathcal{A}, λ) are all equal to 1. Here no CP violation happens and the *mixing induced CP violation* is directly related to the matrix elements of CKM, which are involved in mixing.

Coherent $B^0 - \overline{B}^0$ mixing is mandatory for measuring the mixing induced CP violation in B decays at Belle II, thereby allowing to flavor tag a candidate for B^0 . As a result of the production $e^+e^- \rightarrow \Upsilon(4S) \rightarrow B\overline{B}$, the B mesons are evolving coherently, starting from a well defined quantum entangled state, and are examples of the Einstein-Podolsky-Rosen [14] effect as shown in figure 2.6. The e^+e^- collision at Belle II produces a virtual photon which forms as $\Upsilon(4S)$ due to total energy of the collision being exactly the mass of the B meson. The energy is chosen to be the first bottomonium state above the open B threshold, resulting in decays from $\Upsilon(4S)$ to B mesons in 96 % of all decays.

In order to tag the flavor, one of the two B mesons must decay into its final state. Figure 2.7 illustrates such a decay. The tagged B meson decays via semileptonic decay. Due to the quantum entangled B system, the other B meson from



Figure 2.6: Production of B mesons at the SuperKEKB accelerator via a virtual photon and the bottomonium state $\Upsilon(4S).$

the production via $\Upsilon(4S)$ must be the anti-particle at the time of the decay. The shown decay $B^0 \to D^- l^+ \nu_l$ is flavor specific only for the B^0 . While the other B meson continues to propagate, its flavor oscillates until it decays either as B^0 or \overline{B}^0 . A CP violation can be measured in interference between decay and mixing if the decay happens to be an eigenstate of CP, as in the example of $J/\Psi K_S^0$. The article by Bigi and Sanda [15] concluded that the decay is an excellent candidate for measuring time-dependent CP violations for the β angle (shown in figure 2.3), as it is often referred to as the *golden channel* (see figure 2.8). It is crucial to know the difference in decay times Δt between the two decays of the B mesons in order to determine time-dependent CP violations. It is important to note that in the production process, the B mesons are produced at rest, provided that the $\Upsilon(4S)$ has momentum zero. In this case it is not possible to measure a decay time difference.



Figure 2.7: Schematic of the flavor tagging decay process. In the boosted e^+e^- rest frame an $\Upsilon(4S)$ is produced and decays immediately into two *B* mesons. With one *B* meson decaying via its flavor specific channel, the other *B* meson decays later into a CP eigenstate.

In order to produce $\Upsilon(4S)$ with a non-zero momentum, the kinetic energies of the electron and positron beams must be asymmetric. As show in figure 2.7 by



Figure 2.8: Decay of $\overline B{}^0 \to J/\Psi K^0_S,$ referred to as golden channel.

the different length of the dashed arrows, this mechanism is also used by Belle and Belle II. The electrons are stronger accelerated, such that one can calculate the time difference by

$$\Delta t = \frac{\Delta z}{\beta \gamma c}$$

Here $\beta\gamma$ is the boost of the center of mass frame, c the speed of light and Δz the vertex shift between the decay of two B mesons in beam direction. With the upgrade of the Belle experiment the silicon-strip vertex detector is replaced in the most inner layer by a pixel detector. This decreases the vertex resolution in z direction from roughly 100 µm down to 20 µm. Further details of the Belle II experiment will be discussed later in chapter 3.

The first measurement of the *golden channel* by Belle and BaBar collaborations were published in 2001 [16, 17]. The measured time-dependent CP violation in the *B* meson system is based on the mixing-asymmetry parameter

$$\frac{q}{p} = \frac{V_{td}V_{tb}}{V_{td}^*V_{tb}} \quad . \tag{2.22}$$

With equation (2.18) the connection to the coefficients q and p can be established. With the dominant decay to $J/\Psi K_S^0$ the decay amplitudes can be written as

$$\frac{A_f}{A_f} = \eta_f \frac{V_{cb} V_{cs}^*}{V_{cb}^* V_{cs}} \frac{V_{cs} V_{cd}^*}{V_{cs}^* V_{cd}} \quad , \tag{2.23}$$

where η_f is the CP eigenvalue of the final state. For the decay $J/\Psi K^0_S$ the eigenvalue is -1, but for $J/\Psi K^0_L$ it is +1. The mixing parameter λ can not be written as

$$\lambda = \eta_f \frac{V_{td} V_{tb}^*}{V_{td}^* V_{tb}} \frac{V_{cb} V_{cd}^*}{V_{cb}^* V_{cd}}$$
(2.24)

using the definitions of the angles of the unitarity triangle (equation (2.12))

$$=\eta_f e^{-2i\beta} = \eta_f \cos(2\beta) - i\eta_f \sin(2\beta) \quad . \tag{2.25}$$

The asymmetry amplitudes can be calculated to be

$$C = \frac{1-|\lambda|^2}{1+|\lambda|^2} = 0 \qquad \qquad \text{and} \qquad \qquad S = \eta_f \sin(2\beta) \quad,$$

where the term of the direct CP violation C vanishes and the mixing term S remains. With the previous definition of the asymmetry in equation (2.21), one now can express the time-depended CP violation as

$$\mathcal{A}(\Delta t) = \eta_f \sin(2\beta) \sin(\Delta m_d \Delta t) \quad . \tag{2.26}$$

In equation (2.26) the amplitude $\mathcal{A}(\Delta t)$ and the second sine term can be measured and η_f is defined by the final state of the kaon. The first sine term and therefore also the β angle of the unitarity triangle can be determined. In figure 2.9 the analysis of the full Belle data set of 711 fb⁻¹ is shown.

Over the last decades more results have been published and the CKM-fitter group provides fitted values. In figure 2.10 the latest results from 2021 are shown. The most recent values of the Wolfenstein parametrization [18] are listed below:

$\mathcal{A} = 0.8132^{+0.0119}_{-0.0060}$	$\lambda = 0.22500^{+0.00024}_{-0.00022}$
$\overline{\rho} = 0.1566^{+0.0085}_{-0.0048}$	$\overline{\eta} = 0.3475 {}^{+0.0118}_{-0.0054}$

The angles of the unitarity triangle are

$$\alpha = 91.98^{\circ} {}^{+0.82}_{-1.40} \qquad \beta = 22.42^{\circ} {}^{+0.64}_{-0.37} \qquad \gamma = 65.5^{\circ} {}^{+1.3}_{-1.2}$$



Figure 2.9: Measurements of time-dependent CP asymmetry with the full Belle data set [19]. The left plot shows the CP-odd final state via decay $B \to K_S^0 J/\Psi$. The right plot shows the CP-even final state via $B \to K_L^0 J/\Psi$. The upper plot shows the decay rates in each case, with $B_{tag} = B^0$ in red and $B_{tag} = \overline{B}^0$ in blue. In the lower plots the asymmetry is calculated, which allows the extraction of $\sin(2\beta)$.



Figure 2.10: The global CKM fit in the large $(\overline{\rho}, \overline{\eta})$ plane as a graphical representation of the unitarity triangle. The color areas indicate the error bars. Plot provided and updated by the CKM-fitter group [20].

Chapter 3

The Belle II Experiment

The following chapter describes the basics of the Belle II experiment, its attached accelerator facilities and gives an overview of the single components of the detector. Further the pixel detector is described in more detail with focus on the data acquisition system.

3.1 The SuperKEKB Accelerator

SuperKEKB is an upgraded version of the former KEKB collider that provided beams for the Belle experiment. As the name indicates, SuperKEKB is located in KEKB's same tunnel and uses many of its old components. The principal enhancement, compared to its predecessor, is the increase by 30 in design luminosity. In this way, Belle II will be able to collect data of approximately 50 ab^{-1} . This section provides a brief overview of the main changes required to meet upgrade specifications. A more detailed description can be found in the technical design report [21].

Figure 3.1 illustrates the SuperKEKB accelerator. It is an asymmetric collider and its rings are referred to as the Low-Energy Ring for positrons (LER) and the High-Energy Ring for electrons (HER), each with a circumference of 3016 m. The different energies for electrons (7 GeV) and positrons (3.6 GeV) result in a boosted center of mass system in order to identify displaced vertices of *B* mesons. The boost is decreased from $\beta \gamma = 0.48$ down to $\beta \gamma = 0.28$ due to higher beam currents. This change requires compensation by an upgraded vertex detector scheme of Belle II using a combination of silicon strips and pixels.

A damping ring is used to increase the emittance of positrons by a factor of 2, as their emittance is not small enough to satisfy the injection requirement. As part of the production process, a radio frequency (RF) gun is used to irradiate a photo cathode with laser photons pulsed at short intervals. After acceleration to 7 GeV in the linear accelerator, they are dumped into the HER. The KEKB thermionic RF electron gun is used for positron production. After being fired at a tungsten target at 3.3 GeV and causing bremsstrahlung, an electron-positron pair is produced. As a result of this procedure, a high emitting positron beam is accelerated to a velocity



Figure 3.1: Sketch of the SuperKEKB accelerator facility.

of 1 GeV and directed to a damping ring, which reduces the emittance and further accelerates the positrons to 4 GeV before being stored in the LER.

A significant upgrade to achieve the designed luminosity is the *nano-beam scheme*, which is illustrated in figure 3.2. This technique allows to reduce the vertical beam size at the interaction point by a large amount, thus increasing the luminosity. For two flat beams of equal size, the luminosity is defined in the following manner.

$$\mathcal{L} = \frac{\gamma_{\pm}}{2er_e} \frac{R_L}{R_{\zeta_u}} \frac{I_{\pm}\zeta_{y,\pm}}{\beta_{y,\pm}^*}$$
(3.1)

Here γ is the Lorentz factor, e the electron charge and r_e the classical electron radius. R_L and R_{ζ_y} are reduction factors for the luminosity. They depend on the crossing angle and are of the same magnitude. Adjustable parameters are the beam current I, the vertical beta function β_u^* at the interaction point and the vertical beam-beam



Figure 3.2: An illustration of the collision of beams in the nano-beam scheme. Only the overlap region d is taken into account instead of the much larger bunch length σ_z .

parameter $\zeta_{n,+}$. The beta function is directly connected to the vertical beam size by

$$\sigma_y = \sqrt{\epsilon \beta_y^*} \tag{3.2}$$

Using quadrupole magnets along the beam pipe, the beam is squeezed, and a final focusing magnet (QCS) reduces its vertical beta function to $\beta_y^* = 270 \,\mu\text{m}$ for the LER and $\beta_y^* = 410 \,\mu\text{m}$ for the HER, thereby reducing the vertical beam size significantly. With this method, luminosity cannot be increased arbitrarily and is always constrained by bunch length σ_z , thereby introducing a hard limit of 5 mm in beam direction. By using a finite beam crossing angle 2ϕ , the nano-beam scheme corrects for this effect, resulting in a decrease in the effective cross section of the beams of $d = \sigma_x/\phi$ with σ_x representing the horizontal beam spread. The beta function is constrained by

$$\beta_u^* > d$$

With a half crossing angle of $\phi = 41.5$ mrad and a horizontal beam size of $\sigma_x = 7.75 \,\mu\text{m}$, the effective crossing length is $d = 200 \,\mu\text{m}$, which corresponds to a 25-fold smaller constraint than the 5 mm provided by the bunch length. SuperKEKB will be able to reach its design luminosity goal with the nano-beam scheme of

$$\mathcal{L} = 6 imes 10^{35} \, \mathrm{cm}^{-2} \, \mathrm{s}^{-1}$$

This instantaneous luminosity will be 30 times higher than the Belle luminosity record from 2009 of 2.11×10^{34} cm⁻² s⁻¹ [22] and the total integrated luminosity is expected to reach 50 ab^{-1} by a similar run time. The most recent SuperKEKB luminosity record of 4.71×10^{34} cm⁻² s⁻¹ was reached in June, 2022 [23].

3.2 The Belle II Detector

The Belle II detector is placed around the crossing of the two beams, where particles interact with each other. The detector is a construct of multiple layers housing different subdetectors in order to determine energy, position and time. Each subdetector



Figure 3.3: Sketch of the Belle II detector. Each subdetector is colorized differently. From the outmost KLM in green, followed by the solenoid coil and the ECL in white, the ARICH in blue tiles in the forward region, the TOP as transparent slabs, the CDC as grey long strips and the VXD, which is in yellow for the SVD and in red for the PXD subsystem.

has its own purpose and is placed in a barrel shape around the interaction point or at an end-cap to maximize the acceptance. Acceptance ranges from 17° to 150° in the azimuthal plane and the entire angular range is covered in the polar plane. The higher acceptance in the forward direction is necessary, because particles are boosted into this direction by different beam momenta. The outer subdetectors are responsible for particle identification, the inner ones mainly focus on track reconstruction. In the following a summary of each subdetector is given starting from the most outer ones.

3.2.1 The K_L^0 and Muon (KLM) Detector

The outermost detector focuses on the detection of long living particles, which propagate easily through detector material. For *B* decays the neutral kaon K_L^0 is the most reasonable candidate passing next to muons. Without carrying electrical charge there is no reason to keep the detection of those within the homogeneous magnetic field, therefore the K_L^0 and muon detector is the only one of the barrel subdetectors, which is placed outside of the superconducting solenoid, but still inside the magnetic yoke.
There are multiple layers of glass-electrode resistive plate chambers (RPCs) [24], which were already used in Belle, and of iron plates arranged in alternating order. K_L mesons with a minimum momentum of ~ 0.6 GeV/c (or ~ 1.5 GeV/c depending on their polar angle) will most likely interact in the iron plates and create hadronic showers. Due to a high dead time of the RPC, the detection is adversely affected with higher beam background. Especially in the end-cap regions with limited neutron shielding the background is not acceptable anymore and the RPC layers are replaced by organic scintillator strips [25].

The RPC consists of two float glass layers conducting electricity although they are high electrical resistive. Separated by plastic spacers they are applied with high voltage using a carbon-doped paint and the remaining gap is filled with a gas mixture. Passing charged particles ionize molecules and a streamer formation is enforced by the accelerated electrons in the strong electric field. The streamer formation serves as a conducting channel allowing charge to flow from one plate to another. The high resistive glass limits the size of this local discharge, which itself reduced the electric field and counteracts the streamer formation. With external pickup strips, each with a width of 5 cm, the streamer formation can be transmitted via these strips to discriminators and serializers. Since the strips are mounted on the ground plane of both sides each signal can be detected twice, resulting in a redundant detection.

In the end-cap superlayers of organic scintillators strips are used to handle the higher particle flux mainly caused by background events. The scintillation light is transported by wavelength shifting fibers to photomultipliers, where the signal is converted by frontend electronics.

In order to distinguish muons from charged hadrons, the range of a track whose momentum is the same in the KLM is determined. KLM hits that correspond to muon tracks are used to identify muons in the CDC. Observed and predicted muon ranges are compared to estimate muon likelihoods. False positives are mostly misidentified charged pions, while K_L^0 are identified by hits in KLM that do not correspond with CDC tracks. Neither ECL nor KLM can initiate hadronic showers. K_L^0 candidates must have at least two clusters projecting roughly to the interaction point, where one of them must be in KLM and the second one either in KLM or ECL. A rough momentum reconstruction for the K_L^0 is possible with information obtained from KLM time-of-flight measurements.

3.2.2 The Electromagnetic Calorimeter (ECL)

As a result of *B* meson decays, approximately 3% of the final state products are neutral particles such as pions, kaons, and photons. These produce a large number of photons with energies ranging from 20 MeV to 4 GeV. Consequently, a highresolution *electromagnetic calorimeter* [26] is an integral part of the Belle II's particle identification. The majority of ECL components have been reused from Belle, while only the readout system has undergone significant changes. For the detection, scintillating cesium-iodide crystals doped with thallium CsI(TI) are used. They are read out by photodiodes and digitized by flash ADCs as part of the detection process. For the end-caps, pure CsI crystals will be installed to overcome the more severe background environment. In Belle II, ECL is constructed of 8736 crystals, 1216 of which are in the forward region and 1040 in the backward region. In addition to its primary task, the detection and differentiation of electrons and photons, it also serves as the primary trigger source for Belle II. The energy resolution is expected to be

$$\frac{\sigma_E}{E} = \sqrt{\left(\frac{0.066\,\%}{E/{\rm GeV}}\right)^2 + \left(\frac{0.81\,\%}{\sqrt[4]{E/{\rm GeV}}}\right)^2 + (1.34\,\%)^2}$$

3.2.3 The Aerogel Ring-Imaging Cherenkov (ARICH) Detector

On the forward end-cap inside the ECL is located the Aerogel ring-imaging Cherenkov detector [27]. A 4 cm thick Aerogel radiator is arranged in a radial structure divided into hexagonal tiles at a distance of 167 cm from the interaction point, covering an angle of 14° to 43°. The radiator consists of 2 pads of the same thickness but with different refractive indices (1.055 and 1.065). This way photons produced in the second pad have a larger Cherenkov angle and exactly overlap with in the detector, which provides a better signal yield. 540 hybrid avalanche photo detectors (HAPDs), each divided into 12 × 12 matrices and placed 20 cm from the outer Aerogel layer, detect the light produced.

3.2.4 The Time of Propagation (TOP) Detector

For the barrel, particle identification is done by an *imaging time of flight* detector [28]. It consists of 16 quartz radiator bars of size 2.7 m by 0.45 m with a thickness of 20 mm. A prism is positioned at the reflective edges of each quartz bar to increase the detectable area to 51 mm, as shown in figure 3.4. All borders of these bars, with the exception of one, completely reflect the Cherenkov light emitted by particles traversing the quartz. For detection, 32 micro-channel plate photomultiplier tubes (MCP-PMTs) are arranged in two rows at the prism. Cherenkov light is reflected towards the detection end by a mirror attached to the opposite side of the quartz. By measuring the time between bunch crossing and Cherenkov light detection with the MCP-PMTs, 50 ps time resolution is possible. As a result, the particle travel time to the detector is added together with the propagation time of Cherenkov light in quartz to determine this time. It is possible to calculate the position and angle of the particle when it hits the TOP-quartz by comparing the Cherenkov information with the tracking information. This outcome is compared with a predicted particle, based on the premise that the particle is either a kaon or a pion, which provides the probability of the measured particle being found. It is calculated and associated with the particle candidate in the final reconstruction.



Figure 3.4: Schematic overview and the concept of the TOP counter. A quartz radiator propagates Cherenkov light, which is detected at the bar end by a photon detector.

3.2.5 The Central Drift Chamber (CDC)

As the outermost tracking detector, the Central Drift Chamber [29] consists of 14366 sense wires divided into eight super layers that alternate between each other. There are six sense wire layers in each super layer (the innermost layer has eight) arranged in axial layers parallel to the beam pipe, and stereo layers tilted between -74° and 70° . In the CDC volume that extends from $160\,\mathrm{cm}$ to $1130\,\mathrm{cm}$ in radial direction and 2.4 m in beam direction, helium and ethane are mixed. By traversing the CDC with a curved trajectory, charged particles ionize this gas. The electric field between the wires causes liberated electrons to drift towards the sense wires, which produces an electrical signal. The distance between the origin and the particle trajectory can be calculated based on the signal timing and drift velocity. The trajectory can be fitted as a circle in the plane perpendicular to the beam pipe. To obtain the slope of the typical track-helix, stereo wires have been used. Due to the bending radius of charged particles in magnetic fields, the CDC can determine their momentum. Another key value, determined by the CDC information, is the energy loss per length dE/dx. By ionizing the gas, particles deposit a small amount of energy in each layer when passing through the CDC. It is possible to calculate the energy loss based on the mean energy loss along the particle's trajectory. The correlation between momentum and energy loss is used to separate particle types.

3.2.6 The Vertex Detector (VXD)

As mentioned earlier, the inner tracking detector is divided into 2 completely different subdetectors. The output part of the vertex detector uses the same detection technique as in the predecessor experiment with some improvements. This *silicon vertex detector* (SVD) [30] consists of four layers of silicon strip sensors mounted on ladders. The sensors consist of a p-doped strip in beam direction, a n-doped strip perpendicular to it, separated by an n-doped bulk region, which is ionized by traversing charged particles. During ionization of bulk silicon, electron-hole pairs are formed, and electrons and holes drift towards their charged conjugated strip. Those strips are read out by an external trigger signal, which is sent from the global data acquisition system. A special CMOS-based front-end chip APV25 [31] on the sensor amplifies the signals and transmits them to the flash ADC, which is located outside of highly radiative areas. As the analog signals are shaped and digitalized by the flash ADC, they are transmitted over an optical data link to the global data acquisition system.



Figure 3.5: A 3D model of the SVD. In boost direction the sensor have a wedge shape in order to increase the angular coverage to 17°.

In silicon strip detectors, there is a common geometrical problem associated with their configuration. The detector produces ghost hits as a result of the perpendicular arrangement of the strips. Ghost hits are naturally produced when charged particles cross the detector. Therefore, if a certain number of particles cross the detector at the same time, the same amount of p-doped and n-doped strips will send a signal. Since every hit could have been a potential particle, the DAQ must transmit all signals. This leads to higher background for the SVD in regions close to the interaction point (IP). To overcome these huge backgrounds the inner radius only reaches down to 38 mm from the beam axis and then two layers of pixel-based readout are placed. In figure 3.5 a 3D model is shown.

3.3 The Pixel Detector (PXD)

The inner most detector of the Belle II experiment is the *pixel detector*. The Belle II Technical Design report [21] includes an comprehensive report about it. It consists of two layers of sensors, arranged like the SVD, around the interaction point at 14 mm and 22 mm. In total 40 PXD modules are present, illustrated in figure 3.6. Since this thesis deals directly with the data processing of this subdetector, it will be described in more detail. The physical requirements will be discussed, as well as the exact detection mechanism and the readout and processing of the data. Chapter 4 describes the further processing of the detector data. Figure 3.7 shows the PXD system including supporting frame and cables.



Figure 3.6: Rendering of a PXD module (forward direction, outer layer), by the DEPFET collaboration. On the right side attached Kapton cable connects to the outside. The ASICs are placed next to the active sensor matrix.



Figure 3.7: Photograph of the mounted PXD.

Background Exposition

Before making assumptions about the performance and feasibility of the detector, it is necessary to understand the different background processes to which this detector is exposed. These processes fall in two categories [32]:

Beam-induced processes

arise from the two beams independently rather than from collisions between them. It is generally proportional to the beam current, which only increases by a factor of approximately 2.2 compared to Belle. During its turns in the beam line, a Gaussian shaped beam-bunch loses charge. In order to determine a beam lifetime, it is necessary to consider the typical timescales for each process that contributes to charge losses. Such processes cause particles to cross bunch boundaries and collide with beam pipes, which in turn produce particle showers that, when within the range of a detector, produce unwanted signals overlaying the hits from physics events as background. One of the most relevant processes responsible for beam background is the *Touschek effect*, which involves the intra-bunch scattering of two electrons or positrons. This effect increases with lower beam size, which due to the nano-beam scheme is very small.

A further process reducing beam lifetime and contributes to background is *beam-gas scattering*, which occurs when beam particles escape the bunch and interact with residual gas molecules in the beam pipe.

Lastly, *synchrotron radiation* is the result of charged particles being accelerated, causing a photon to be emitted. The electron and positron beams produced at SuperKEKB are deflected by quadrupole magnets to keep them in the circular beam pipe. Magnets before the interaction point focus the beams, which contribute to additional background hits in the detector.

Luminosity-dependent processes

are induced by QED interactions of electrons and positrons, directly associated with beam collisions. The predominant source is the dominant *two-photon process*.

$$e^+e^- \rightarrow e^+e^-\gamma\gamma \rightarrow e^+e^-e^+e^-$$

Since they are directly dependent on instantaneous luminosity, their rate is expected to be 30 times higher than at Belle. During the photon-photon interaction, low energetic electrons and positrons curl in the magnetic field, leaving signals only in the PXD.

Bhabha scattering is another source of background from QED processes. Colliding electrons and positrons do not annihilate but scatter. Afterwards, these scattered particles leave the interaction region under very small angles and are converted into secondary particles in the beam pipe.

Based on a background simulation [32], the occupancy is calculated. It describes the number of hits relative to the total number of PXD pixels in a single event. All background sources add up to 1.74% occupancy. Less than 0.1% of the pixels are triggered by actual signal hits, which are useful for later analysis. In light of the hardware limit of 3%, the PXD should function effectively under the expected conditions. Assuming that the design trigger rate is 30 kHz, the total PXD data rate is around 20 GB/s. This requires a dedicated data acquisition system with fast readout and reduction capabilities to handle background dominated data.

The actual background situation of the full Belle II detector in 2021 is summarized in the Snowmass 2021 paper [33].

Detection Mechanism

For the detection, a depleted field-effect transistor (DEPFET) [34] is used. This detector technology is utilized in Belle II as part of a particle physics experiment for



Figure 3.8: Illustration of a single DEPFET pixel.

the first time. In order to achieve a good vertex resolution, multiple scattering of particles within the detection layer must be kept as small as possible. Consequently, the thickness of the sensitive area was of high importance. The DEPFET technology allows a thickness of only $75 \,\mu$ m.

The figure 3.8 illustrates the main features of a DEPFET pixel that is composed of silicon soils (bulk) grounded on positive contacts and topped with p-channel MOSFETs¹. The positive grounding plate causes the bulk region of the DEPFET to become fully depleted. At the same time the p-doped ground plate acquires a negative potential. An electron-hole pair is created in the silicon soil by ionization of a traversing electron. Positively charged holes drift to the ground plate while electrons accumulate in a potential minimum below the gate. When the gate is active, it is possible to measure the drain current, which is directly proportionate to the electron count. In order to perform the next measurement, the electrons within the internal gate need to be cleared once the drain current has been measured and the amplified signal has been sent down the data acquisition chain. When a positive potential is applied to the clear gate, electrons will drift to that location and allow a new readout cycle to take place.

DEPFET pixels are arranged in an array of 768×250 for a single module. Each pixel measures $55 \,\mu\text{m}$ times $60 \,\mu\text{m}$ in the inner layer and $70 \,\mu\text{m}$ times $80 \,\mu\text{m}$ in the outer layer. With the supporting frame the thickness is increased to $420 \,\mu\text{m}$. All modules have a sensitive area with a width of $12.5 \,\text{mm}$, extended to $15.4 \,\text{mm}$ by the supporting structure. The length varies between the inner and outer layers. While the sensitive area in the inner layer measures $44.8 \,\text{mm}$, the length is slightly higher

¹Metal Oxide Semiconductor Field Effect Transistor

in the outer layer at 61.44 mm. The frame extends each module to 68 mm for the inner layer and 85 mm for the outer layer. The surrounding rim increases mechanical stability and provides space for application-specific integrated circuits (ASICs), which are essential for control and read-out as well as voltage supply. A flexible printed circuit (Kapton) establishes the connection to external control systems and power supplies.

Read-Out

A DEPFET module has 250 long columns and 768 short rows of DEPFET pixels arranged in a matrix. Electrically these pixels are rearranged by combining four rows into a row group, yielding an array of 1000 logical columns, called *drain lines* and 192 logical rows, called *gates*. The modified logical structure results in an increased readout speed. The voltage is applied to the DEPFET in such a way that the accumulated charge can flow to the ADC via drain lines. All logical columns are read out simultaneously, while only one logical row is supplied with the necessary voltage. The logical rows are processed one after the other to ensure readout and complete discharge. This way 192 pixels are connected to the same ADC channel.

The readout of a module requires three different ASICs [35], which are placed directly next to the matrix.

Switchers

The six switches apply the gate and clear voltage. All pixels within a logical row are operated at the same voltage. There are three operating modes. While both voltages are held back, charge can be deposited in the sensor. The module is in this state until the readout of this row is started. Then the gate voltage is applied, so that readout is possible via an ADC. The clear voltage is added with a delay to clear the collected charge. The readout and clear of a row takes about 100 ns before the switcher continues with the next row. With almost 200 rows, the complete readout of a module takes $20 \,\mu\text{s}^2$.

DCDs

4 DCDs (drain current digitizers) [36] convert analog to digital signals. Each DCD has 256 channels with 8 bit resolution. Additional features like individual offset correction and adjustable gain and bandwidth allow individual adaptation to each sensor. Dark currents can be individually subtracted, increasing the dynamic range of the ADC. The load of 1000 logical columns is measured via the drain line on all DCDs. After processing, the data is sent to the respective control center, the DHP, via eight 8-bit links.

DHPs

4 DHPs (data handling processors) [37] store incoming data in a ring buffer. The buffer size is selected so that the data of an entire module can be buffered twice. This ensures that all data from the last complete readout is available even after the trigger signal arrives. DHPs can configure DCDs via JTAG connections. At the same time, one of the DHPs controls the voltage sequencing of all switchers. With the trigger signal, further processing is started and data is forwarded via high-speed optical links. During processing, digital pedestal offsets are subtracted from each pixel value. This offset can be set individually for each pixel and is usually determined using raw data.

 $^{^2} For$ comparison, each bunch in the accelerator passes the interaction point twice during 20 $\mu s.$ Therefore the charge and number of signals are accumulated for multiple events.

Chapter 4

Data Acquisition

Due to the complexity of the Belle II detector, it is not possible to analyze the analog signals on the fly. Instead, the information must be stored for later offline analysis. Therefore, the signals from each detector component must be read out, digitized and stored. The conversion of analog signals to hard disk storage is the purpose of the data acquisition (DAQ) of the Belle II detector. Although the DAQ is available from the previous Belle detector, the level of detail and trigger rate is increased and requires an upgrade of the entire DAQ system. The 40-fold increase in data rate of the predecessor is not being dealt with by the existing implementation and an upgrade in hardware and software is foreseen. Six gigabytes of raw data per second will be produced by the outer detectors, while the high resolution of the PXD will result in more than 15 gigabytes of unfiltered data per second.

The data handling of the outer detectors is unified at each step of the data acquisition chain in order to minimize the maintenance cost to operate the system for more than a decade. The PXD has a separate data acquisition and all raw data is truncated before being written to disk. This behavior is explained in this chapter.

4.1 Belle II DAQ Overview

4.1.1 Trigger System

The data from the Belle II experiment is separated into individual events. Each event is issued by a trigger system linked to the Belle II physics program. Each field of interest can be described by the subdetectors using a variety of characteristics. Configurable logic based on FPGAs allows a fast response after combining raw signals and generating a draft version of an analysis. If the signature of this rough analysis matches any of the predefined characteristics of interest, a single trigger signal is issued.

Many of the pre-defined signatures contain signals from at least two charged particles with their tracks coming from the interaction point. The charge and track direction can be extracted from the CDC and ECL, which serve as the main input for this rough analysis. The calculation of the energy deposition of a cluster in the ECL with its track shape allows to distinguish between different types of electromagnetic processes. Additional timing information from the ARICH and the TOP or the presence of muons from the KLM can provide further information to trigger events. With redundant tracking and calorimeter trigger information, Belle II provides greater than 99 % efficiency for *B* meson events.

All channels send their data to the *global decision logic* (GDL) [38, p. 1], summarizing the available signals and creating a final trigger signal. The GDL configuration allows to mask or pre-emphasis certain channels, if necessary. The final trigger signal is forwarded to the *frontend timing switches* (FTSW) [39], an FPGA based distribution system, which ensures simultaneous readout of the frontend digitization boards of the most recent data. Multiple FTSWs are arranged in a tree-like structure. The main FTSW generates a *level-1* trigger signal that includes precise timing information and an event number that automatically increments with each received signal.

Based on a study of Belle, the level-1 trigger rate is expected to be nearly proportional to the luminosity and is approximately 20 kHz at the design luminosity [40]. With contingency, the entire data acquisition system is designed for a trigger rate of 30 kHz.

4.1.2 Data Flow

When the trigger signal arrives at each subdetector, the readout is started by the fronted digitization boards. The data flow is illustrated in figure 4.1. The detailed stages are described in the following sections. The PXD system has a completely different time and readout structure and its data stream is separated until the last stage.



Figure 4.1: Overview of the Belle II data acquisition system.

FEE

All subdetectors provide dedicated hardware directly connected to their sensors to handle the low level configuration and start the readout process after receiving a trigger signal. This *front-end electronics* (FEE) is based on FPGAs or ASICs to process the trigger signal as fast as possible and to stream the necessary information as raw data to the next level. The analog pulse information is already digitized at this first stage. Different data schemes are configurable to take different data for physics or calibration data acquisition.

COPPER (Digitizer)

The readout process is the same for all detectors. Receive the level-1 trigger, transmit the digital data using a serialized data stream, receive slow control signals such as calibration constants, and receive a system clock to operate the system. This steering mechanism is provided by the *common pipeline platform for electronics readout* (COPPER) module [41]. Each COPPER module consists of 4 daughter cards that receive the data. The COPPER system has been verified to handle up to 40 kHz of level-1 trigger rate. In case of bandwidth limitations, the number of receiving slots per COPPER can be reduced. The event fragments are received from up to 16 COPPER modules and transmitted via Gigabit Ethernet to a *readout PC* (ROPC).

Event Builder (EB) and High Level Trigger (HLT)

The *event builder 1* (EB1) is a computer farm receiving data from all ROPCs. A round robin scheme is used to build up full events for each subsystem, therefore a full mesh topology is used. From here on, the increased event size requires 10 GbE connections. The events are distributed with correct order from all EB1s to the HLT simultaneously.

The *high level trigger* (HLT) [42] is a server farm, where finally 20 units process events in parallel. Each HLT unit consists of a distributor node, multiple worker nodes and a collector node. The HLT performs a fast full reconstruction of all received data in order to make a physics-relevant decision as to whether or not save the data for later physics analyses. The processing time for a single event varies between a few milliseconds and five seconds. In average the processing time is expected to be less than one second. The data is transmitted to the next stage as soon as possible, while the incoming order is not preserved. PXD data is not included in the decision process. In case of storing the data, *regions of interest* (ROI) are generated to truncate the PXD data stream. Details will be discussed in section 4.3.

The HLT can operate in two modes: *Filtering*, when the decision is forwarded to the next level, where the raw data may be dropped, and *Monitoring*, where the decision is always to keep the raw data. The latter mode aims to verify general operation without losing data. It is not foreseen to run this mode at design luminosity, where the data rate is too high.

The *event builder 2* receives the data from the HLT and ONSEN and builds final events before sending them to storage. The data is stored in SROOT files [43], which prevents data corruption in previous events in case of failure. Such SROOT files must be converted to normal ROOT files [44] for physics analysis, which is done offline.

4.2 PXD Data Handling

As shown in figure 4.1 the PXD data stream is separated from the common scheme. Since it takes $20 \,\mu s$ to read the entire PXD, there is a high probability that the next trigger will arrive before a single read cycle is completed. Therefore, the entire detector is always read, while the level-1 trigger signal is used to mark the event fragment to be passed to the next stage. All hits are accumulated over a full readout cycle and will occupy roughly 2% of the pixels. With a maximum average occupancy of 3% and a trigger rate of $30 \,\text{kHz}$ the readout system must be able to handle up to $18 \,\text{GB/s}$.

Data rate
$$\approx \underbrace{7680000}_{\text{pixel}} \cdot \underbrace{0.03}_{\text{avg. occupancy}} \cdot \underbrace{30 \text{ kHz}}_{\text{trigger rate}} \cdot \underbrace{2.5 \frac{\text{bytes}}{\text{pixel}}}_{\text{data size}^1} = 17.3 \frac{\text{GB}}{\text{s}}$$

This data rate is about an order of magnitude higher than the summed data rate of other subdetectors. Therefore the common unified data acquisition system is not capable of handling the PXD data stream. Additionally there is no reduction mechanism foreseen and the raw data would be too much to store permanently. Before the PXD data is merged with the data from the other subdetectors by the event builder 2, it passes through several stages. All processing steps after receiving data from the DHPs will be explained now.

The *data handling hub* (DHH) [45, 46] is referred to as the entire system, which controls and reads data from the PXD modules. It consists of 8 *data handling concentrators* (DHCs), 8 *data handling insulators* (DHIs) and 40 *data handling engines* (DHEs). One DHC, one DHI and 5 DHEs are hosted and connected by a DHH carrier card (DHHCC). The DHE and DHC are designed as single-width advanced mezzanine cards (ACMs), while the DHI requires twice the width. The DHC is connected internally by single high-speed serial links to all other cards hosted by the same DHHCC. External connections are established by optical SFP+ gigabit transceivers for data acquisition and slow control and Infiniband cables to the sensor Kapton of 5 PXD modules via a patch panel. CameraLink cables are attached to the DHI used for configuring, controlling and triggering 5 PXD modules. All DHH modules use FPGAs and are equipped with 4 GB of memory and can handle internal connections up to 5 Gbps. External optical connections allow higher transfer rates of 6.25 Gbps. Figure 4.2 illustrates the internal and external connections.

 $^{^1}Each$ firing pixel results in 2 to 4 bytes of data in zero-suppressed data format. Simulations have shown, that the average pixel size results in 2.5 bytes/pixel. The header sizes can be neglected at 3 % occupancy.



Figure 4.2: Schematic of the connectivity of DHH. The DHC is the central system, which forwards the trigger signal via the DHI to the modules. Afterwards the data is sent from the DHPs mounted on the module to the DHEs and the DHCs. The DHE and DHC perform subevent building. Each DHC sends events to four ONSEN nodes using a round robin scheme.

The DHE is used to receive data from a single PXD module. Here, time information and DHP internal processing must be considered carefully. Due to the long integration time of $20 \,\mu s$ caused by the rolling shutter mechanism, data receiving must be started with an incoming trigger signal. This is executed for one readout cycle, so that exactly 1 image of all pixels of a module can be stored. If another trigger signal is received within this process, the data must be stored in several events from this point on. For this purpose, the DHP behavior is emulated in the DHE.

The DHC receives data from 5 DHEs and builds subevents from them. Afterwards, the events are sent to ONSEN. An event is only sent via one of the 4 optical AURORA links to balance load. Part of the data is used for module monitoring. Thereby the data of the first of the 4 links is additionally delivered via UDP to a local DAQ system, which analyzes it independently of Belle II's data quality monitoring (DQM). For detector studies it is possible to send all data via this output, e.g. for module calibration.

The DHI controls 5 PXD modules via CameraLink cables. A clock signal, a control signal and JTAG commands are transmitted. For the control signal, 4 individual functions are possible, sent continuously: reset, veto, trigger and synchronization. Via JTAG the DHP configuration can be adjusted and read back. However, due to the permanent JTAG connection, the optical transmitters of the DHPs are supplied with 3.3 V, so that the DHPs do not completely lose their configuration when their operating voltages are switched off. For this purpose it is possible for the DHI to cut

this connection to reset the DHP configuration completely².

When the DHCs receive a trigger signal from the Belle II Trigger and Time Distribution System (B2TT), they forward the signal to the DHEs and DHIs. The DHIs henceforth send a trigger signal to the DHPs for $20 \,\mu$ s, so that the readout from its memory begins. Since the trigger signal from the B2TT is independent of the rolling shutter readout position, data is sent to the DHE over up to 2 frames. Since only 1.5 Gbps are available for each optical link from the DHP, the average occupancy rate per module of 5 % may only be exceeded for a short time. This is to prevent data loss. The DHEs receive the data and emulate a possible end using the DHC trigger information. This is necessary because DHPs only send hits above a certain threshold after the first processing. After that, the DHEs send the events to the DHCs, which assemble the data from 5 DHEs into subevents and forward them. The maximum outgoing data rate is $625 \,\text{MB/s}$. Due to the different expected data rates in inner and outer PXD modules, data is always received from 2 inner and 3 outer PXD modules per DHC. Figure 4.2 shows the trigger and data flow.

4.3 Online Selection Node

In this section, the requirements for handling up to $18 \, \text{GB/s}$ of raw PXD data are explained. Afterwards the hardware and operation of this system are explained in more detail.

As with the outer detectors, the first step to PXD data reduction is the HLT event decision by a factor of 3. PXD data aren't processed by the unified data-acquisition scheme, so there must be a different data path set up. Since the PXD modules are triggered by the level-1 trigger, the data must be buffered until the HLT decision arrives. A dedicated system must wait and send only the accepted events to EB2 for integration with other subdetectors' data. As a first requirement, a PXD data reduction system must have a large input bandwidth compatible with DHH outputs of 6.25 Gbps. It must also have an output interface to the EB2, and sufficient memory capacity and bandwidth. The necessary memory can be estimated to be 36 GB using twice the average HLT processing time of 1 second and the previously calculated data rate.

Due to the background domination in the PXD, its data consists mainly of not meaningful information. Analysis of the pure PXD data is ineffective and requires particle tracks found in CDC and SVD, where a much lower background level allows reconstruction without an exorbitant number of fake tracks. The combination of such trajectories with PXD information allows improved vertex resolution around the interaction point.

The data reduction process is logically extended to exclude PXD hits that cannot be associated with SVD or CDC tracks. The HLT implements this algorithm as part of its online event reconstruction algorithm. Based on hits in SVD and CDC,

²The JTAG communication of the DHI to the PXD modules used to be handled by the DHE. But due to the back powering of the DHP, additional hardware was necessary that controls this problematic part via slow control.



Figure 4.3: Illustration of ROI calculation. A charged particle moves outward on a helical trajectory from the interaction point on the active surface of the PXD and half of the SVD layer, leaving hits in each strip of the SVD. As a result, a particle's track can be reconstructed, and the two PXD layers are extrapolated to determine the most likely intercept locations. Regions of interest around those intercept locations are defined.

particle tracks are extrapolated back to the interaction region, estimating the vertex position. In the area around the interaction point, particles will produce PXD hits. Therefore, rectangular areas in both layers of PXD are marked as a *region of interest* (ROI). Figure 4.3 illustrates the basic concept. PXD data reduction is based on ROI information, along with HLT decisions.

The rejection of all PXD hits that do not fall within ROIs eliminates the majority of background data and leaves only a small amount of signal-related data. As a result, it is possible to guarantee that the ROIs will include all pixels relevant to the PXD intercept based on the resolution of the track position at the point of the PXD intercept. At this stage, the event-building system requires a reduction factor of 10, which means ROIs cover approximately one tenth of the entire sensor area. A total data reduction factor of 30 is achieved by combining both concepts. This reduces the PXD data rate to 600 MB/s.

Every event triggered by a level-1 trigger is also accompanied by a second set of ROIs produced by the *Data Acquisition Tracking and Concentrator Online Node* (DATCON) [47, 48]. DATCON is an FPGA-based system that utilizes the same hardware platform as the PXD data-reduction system described later in this section. It exclusively utilizes SVD data. Similar to the HLT concept, it reconstructs SVD tracks, extrapolates the tracks to the PXD layers, and defines ROIs around the intercepts. For each event, the output from DATCON will be available significantly faster than from HLT and must be buffered. In addition to the redundant nature of the two ROI sources and the need to logically combine the selected areas, only the HLT can select the events for storage. As a result, DATCON ROIs for events that do not qualify for event selection by the HLT are not taken into account.

In addition to providing interfaces to the HLT and DATCON, the ROI mechanism adds additional prerequisites to the data reduction system. Additionally, the system must be able to perform matching of ROIs from both sources to the correct data and filter the hits accordingly.

4.3.1 Hardware

With the requirements described in the beginning of this section, the *online selection node* (ONSEN) was developed [49, 50, 51]. The ONSEN system consists of 9 *compute node carrier boards* (CNCBs) and 33 *xTCA-based FPGA Processor* (xFP) cards. The design of all boards was developed in cooperation of the Institute of High Energy Physics (IHEP) in Beijing, China, and the Justus Liebig University (JLU) in Gießen, Germany. All design choices were made following telecommunication standards defined by the PCI Industrial Computer Manufacturers Group (PICMG) consortium [52, 53, 54].

The ONSEN system is hosted inside an *Advanced Telecommunications Computing Architecture* (ATCA) shelf with 14 slots. It provides a full mesh topology, which allows each daughter carrier board to communicate with each other at data rates up to 3.125 Gbps. It supports the *Intelligent Platform Management Interface* (IPMI) standard for monitoring and controlling the whole shelf and daughter boards. Via two redundant shelf managers accessible via TCP significant information is provided by the slow control system, which allows remote access and monitoring of power consumption and temperature.

The CNCBs are ATCA carrier boards installed in 9 of the 14 slots on the ATCA shelf. Each CNCB hosts a Xilinx Virtex-4 FPGA [55] and a 2 GB DDR2 SDRAM module. Additional components like serial hubs and an attached IPM controller (IPMC) are placed on the PCB as well. Multi-gigabit connections to the back plane of the shelf and low voltage differential signal (LVDS) links to the daughter xFP cards are routed to the FPGA. On the back side it can be extended with a *rear transition module* (RTM), which gains access to an Ethernet port, a USB hub and the JTAG chain. The RTM Ethernet port provides a slow control interface. To perform debugging, the USB hub facilitates universal asynchronous receiver/transmitter (UART) communication. The JTAG connection allows firmware uploading. For automatic programming a 64 MiB flash chip and a complex programmable logic device (CPLD) is added to automatically upload the firmware to the FPGA. A combined bitstream is used, which holds the carrier and xFP cards' bitstreams. The flash also holds the compressed Linux for the carrier, which is executed by the PowerPC after firmware upload. Figure 4.4 shows a picture of a single CNCB.

The smaller xFP cards are advanced mezzanine cards (AMC) with a Xilinx Virtex-5 FPGA [56] and two 2 GB DDR2 SDRAM modules. Two multi gigabit connections from



Figure 4.4: Photograph of the CNCB with 4 xFP cards installed. The plate on top ensure air flow when installed in the ATCA shelf.



Figure 4.5: Photograph of the xFP AMC card equipped with a cooling block on top the FPGA and two DDR2 SO-DIMM modules.

the FPGA to SFP+ transceivers allow external optical connections up to 6.125 Gbps. Via 4 LVDS links to the carrier's FPGA a combined data rate of 600 Mbps can be achieved. An additional Ethernet port provides an interface to external monitoring. UART connections to the front and the AMC connector provide debugging functionality. Sensor readout is achieved via a *Module Management Controller* (MMC), which communicates with the IPMC on the carrier. A flash chip holds the Linux for PowerPC of the FPGA. A picture of a single xFP card is given in figure 4.5.

4.3.2 Firmware

There are four different firmware versions available for ONSEN. One of the 33 xFP cards is the *Merger* node. It receives data on optical links from DATCON and the HLT. DATCON data is sent via a 3.125 Gbps AURORA link and stored in memory. After event processing the ROI packages from HLT are collected by the ROISENDER node and forwarded via a Gigabit Ethernet link. With the data from DATCON being already present at the arrival time of HLT information, the merging of HLT and DATCON ROIs is executed. For later filtering the DHE-ID must remain in ascending order during the merge procedure. Afterwards the merged events are transmitted to the host CNCB's FPGA for further distribution.

The other xFP cards are called *Selector* nodes. They receive the pixel data from 5



Figure 4.6: Schematic of the connectivity of ONSEN. Each HLT unit receives detector data from EB1 and calculates ROIs for each event. At first the ROI information is sent to the Merger, where ROIs are combined with DATCON ones. Afterwards the merged ROIs are distributed to each Selector as indicated by the dashed lines. After processing the DHC data, the Selector sends the reduced pixel data to EB2, where it is merged with the rest of the data before being stored.

PXD modules and buffer it in the memory until the merged ROIs from their CNCB's FPGAs arrive. The buffering mechanism is similar to the one in the Merger node, where the trigger number and the memory location are stored in a look-up table (LUT) and later based on the trigger number of the HLT data the correct data is read again. In accordance with the ROI and the HLT decision, single hits are retained or discarded. Afterwards the processed data is sent to EB2 over a Gigabit Ethernet link.

The CNCB FPGA firmware differs slightly between hosting the Merger node and Selector nodes. The *Merger Switch* distributes ROI data to 8 Selector Switches. Due to the load-balancing mechanism on the DHC, only every fourth event is sent to each Selector Switch. This way all Selector nodes of one CNCB process the same event. The *Selector Switch* distributes ROI data to 4 Selector nodes. Since each Selector node only receives pixel data from up to 5 PXD modules unused ROIs with different DHE-ID are dropped.

Figure 4.6 illustrates the four ONSEN subsystems on the right, along with external data sources. The dashed lines show the flow of ROI information, whereas solid lines show the flow of detection data.

Most ONSEN firmware IP cores provide one or more 32-bit interfaces to pass payload data. They utilize the Xilinx LocalLink protocol [57], which offers a simple flow control mechanism to prevent data loss between different cores. These internal connections do not verify the data by a cyclic redundancy check (CRC). This mechanism is only used for external connections using the AURORA or TCP protocol. If verification fails, the entire event is dropped.

During operation, it is essential to have a controlling and monitoring mechanism. It is possible to accomplish this by using firmware designs that include CPUs, also known as *embedded systems*. The Xilinx Embedded Development Kit [58] helps achieve this. For ONSEN the firmware is built around the PowerPC available in the Virtex-4 and Virtex-5 FPGA. Distinct firmware design blocks called intellectual property (IP) cores are connected via a local bus [59] to the CPU. Communication with each core is made possible by reading and writing to the bus. The most basic feature is to expose 32-bit slave registers, which can be used by the firmware to provide status information or by the CPU to apply configuration during operation. The IP core can send interrupts to the CPU, which can efficiently wait and dispatch a notification.

Software running on the embedded CPU is supported by the Xilinx *Software Development Kit* (SDK). Standalone programs or more complex operating systems (OS) may be used. However, there is only a limited amount of space available inside the bitstream. For the PXD, the *Experimental Physics and Industrial Control System* (EPICS) for slow control has been chosen. Since the EPICS ecosystem depends on various Linux libraries, the underlying OS had to be prepared using an open source build system called Buildroot [60]. It allows to generate a Linux kernel by cross-compilation for a different architecture. A device tree extracted from the firmware design includes the necessary information to map each core to a memory address accessible from the OS via the local bus.

The OS and software are finally packaged into an *Executable and Linkable Format* (ELF) [61] file and executed by a small bootloader, which is included in the bitstream. The bootloader decompresses the ELF file, which is typically stored on the permanent flash chip. It boots until all the desired software is operational. 256 MiB of the DDR2 memory is reserved for the OS. Further details on the software will be discussed in the next section.

In the following contributions from the work on this thesis will be explained.

PXD Data Verification

Receiving data from external sources uses the AURORA or TCP protocol. While the TCP protocol verifies the integrity of the data by a CRC and supports retransmission, the AURORA protocol does not offer these functionalities. Retransmission of AURORA frames is not included in its specification and would require extensive work on the sender and receiver sides. Data integrity can be assured manually. It is imperative to know that the PXD data format is divided into ROI and pixel data. A single ROI event consists of a single frame with 32 bytes of meta data, a multiple of 8 bytes of payload data and a final 4 byte checksum. Since each processing instance of the DHH, i.e. DHE and DHC, surrounds the current event by a start and end frame with meta information and CRC, a single pixel event contains multiple frames. The DHP data frame is the only pixel data frame that contains payload data. The different frame types are described in detail in the appendix of the thesis by T. Geßler [62].

For this purpose the *Belle II Format Handler* core connects to the output of each AURORA core. While the data passes through, the CRC is calculated and matched against the last word of each frame. Additional verification procedures are applied for each frame type, e.g. whether additional meta information is set correctly or the frame size of each DHH frame is correct. On top, the data is analyzed on the fly and

consistency is validated.

If an event contains at least one error, the whole event is discarded. In the case of the ROI data from DATCON and the pixel data from the DHH, the data is not retained in a buffer, but directly passed to memory. In order not to use this event for further processing, the stored event is flagged as faulty and no entry in the look-up table (LUT) is added. The data may remain in memory for debugging. The ROI data from HLT or already merged with DATCON is buffered until the CRC is verified.

During the work on this thesis, the checks have been extended and made optional. Many of these errors were seen in offline data analysis and pointed to DHH firmware issues. Error checks have been implemented at the firmware level to detect them as soon as possible without losing valuable time. There are 13 mandatory checks for pixel data and 14 optional checks that can be configured via slow control when it comes to pixel data. New error checks are described below.

DHH data counter

The DHC and DHE store the number of enclosed data in their final frame. This information is later used to draw conclusions about the data reduction after filtering by ONSEN. The DHC stores the number of 32-bit words in a single 32-bit word and the DHE counts the number of 16-bit words. Since the Belle II Format Handler has to measure the pixel event size for storage in memory, the relevant counting mechanism is already present.

Start with Double-Row

The zero-suppressed (ZS) data frame from the DHP was previously forwarded by the DHE. ZS payload data consists of a 16-bit word for each double row followed by one or more 16-bit words for each hit within this double row. To calculate the exact location of each hit, the double row information must be given before the hit information. With the DHE firmware upgrade to handle overlapping triggers³, DHP data is processed by the DHE. In some situations this led to missing double row information in the first payload data word. In this case the following hits cannot be reconstructed and the information is useless. During development of the DHE firmware it also happened that multiple double row words appeared at the beginning, which do not make sense. The first two 16-bit words are confirmed to be a double row word followed by a hit word.

Consistent DHH ID

The DHC and DHE opening and closing frames contain their identities. In some rare cases the subevent building of the DHC did not work as expected and the DHE data was scrambled. Afterwards the pixel information cannot be correctly associated with its attached PXD module and the entire event is worthless.

³Overlapping triggers happen, when the next trigger signal is received before the readout of the current event is finished. In such case the DHE duplicates the data for the next trigger.

Ascending DHE ID

The DHE ID in pixel and ROI data must be ordered ascending for the filter mechanism of ONSEN. A check on the correct order of each pixel event is added.

For the extraction of the pixel count, which will be explained later, it is also essential that the order of the DHEs remains the same, i.e. data from a single DHE must not be missing for individual events.

Consistent Trigger Number

The DHC and DHE opening frames and the DHP frames contain the trigger number received by the trigger system. To prevent scrambled events, the trigger number in the DHC opening frame is compared against the value in all other frames. Later this check was enhanced to differentiate between the consistency of the lower and upper 16 bits of the trigger number, to narrow down the source of the problem.

With the Belle II Format Handler analyzing the outgoing data stream to the event builder 2, the trigger number is also compared against the trigger frame⁴ and the optional ROI frame.

PXD Data Analysis

With the full processing of the pixel data already happening in the Belle II Format Handler core, its functionality was enhanced to extract certain information from the data. This information is extracted into slave registers read out by the PowerPC periodically and published as EPICS PVs.

In addition to the word counts, the hits within the ZS data frames are counted for up to 5 modules individually. The DHE ID is extracted for the first event and the number of all pixels is added. Due to missing support for unsigned integer values in EPICS, only 31 bits of the 32-bit counter are forwarded to the slave register. Finally the *average occupancy* is calculated in EPICS, with the difference between the current and previous values and the event number difference. The average occupancy serves as the highest quality feedback on how each module behaves under current conditions. This is so that beam parameter changes can be seen immediately. Figure 5.3 shows the calculation from values of 4 selector nodes.

However, the mean calculation has the disadvantage that single events with high occupancy, so called *burst events*, are not visible especially at high trigger rates. Therefore the core was improved to store the maximum number of pixels within an event as well. The value is read out via EPICS and a reset is applied to the firmware. This way the *maximum occupancy* within the typical time window of 0.5 s can be calculated.

If too many burst events happen in a row the DHP cannot send all data. The average occupancy limit stands at $2.5\,\%$, while the maximum occupancy is $30\,\%$ for

⁴The trigger frame is added at the beginning of the outgoing data, including HLT meta information. This is to properly match HLT data with ONSEN data.

single events. If data is truncated the common mode value in the data stream is set to 63, its highest value. The firmware also counts the number of events where such a value is seen. This allows to see data truncation on the fly. With this information being directly related to noisy injections, these values are visible in the accelerator control room.

The consistency check on the trigger number offers the possibility to pose the current trigger via EPICS. This is done on the merger and selector nodes. Using the difference between these values and the trigger rate, the delay between level 1 trigger and HLT processing can be calculated⁵. In this way, it is possible to verify that the maximum delay of 5 seconds has been met.

4.3.3 Software

Several applications run under Linux on the PowerPC. The most extensive part is the EPICS system for slow control, which is explained later in section 5.2. Other applications help manage the system and reduce manual initialization.

The applications are built via a cross-compiler, part of Buildroot, and packaged into the bootable Linux container. With limited flash chip capacity, individual components must be small. This leads to a reduced number of libraries available for custom applications.

Additional applications and configuration files can be placed in a separate part of the flash, which is mounted as a filesystem under Linux. But this approach has the disadvantage that the file version cannot be guaranteed easily, so packaging binary files is more efficient.

Apart from EPICS, a low-level application monitors all slave registers and interrupts. Additionally it handles initialization procedures for all cores, initializes the LUT and resets the memory. Based on a configuration file on the mounted flash, it configures the IP address of the SiTCP core. This core establishes the connection to EB2. The application connects to EPICS, allowing it to be controlled by the run control. This application's output as shown in figure 4.7 looks very minimalistic because it is based on a version without an operating system.

IPMI Communication

ONSEN nodes have an Ethernet interface to transmit slow control information. All nodes operate on the same subnet and therefore require a different network address. The IP address and hostname were set in a configuration file on the flash chip. However this requires manual changes in case the CNCB or xFP location is changed due to replacement or debugging.

Since the PowerPC has access to a serial connection to the IPMC or MMC, it is possible to request the exact location inside the ATCA shelf. During boot the address

 $^{^5 {\}rm The}$ difference does not only include the pure processing by HLT, but also the processing in event builder 1 and all data transmission.

. 013S1 T: 67 degC U_1V0: 984 mV U_2V5: 2499 mV									
[l]oad [u]nload [s]tart sto[p] [a]bort [d]ump [w]rite lu[t] [q]uit									
<aurora_amc> - v1.06.a</aurora_amc>									
AUS	00010001	00010000	fffffff	ffffffe	00000001	00010000	I:1		
AUM	00020002	00020000	fffdffff	fffdfffd	00020002	00010000			
BFF	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	
BPR	00000000	00000000	785ffdc0	005ffe00	0017fff3	00000000	00000000	00000000	
FMT	0000000a	00008031	00000000	fff80000	00000000	00000000	00000000	00000000	I:0
	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	
	00000000								
FMT	0000000a	00000030	00000000	fff80000	00000000	00000000	00000000	00000000	I:0
	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	
	00000000								
FMT	0000000a	00000031	00000000	ffe00000	00000000	00000000	00000000	00000000	I:0
	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	
	00000000								
MON	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	
NRD	00000002	0000003	00000000	I:0					
NRD	00000003	00000003	00000000	I:0					
NWR	00000002	0000003	I:0						
NWR	00000002	0000003	I:0						
PXF	0000002a	00000007	I:0						
LRD	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	I:0
LWR	00000000	00000000	00000000	00000000	00000000	00000000	00000000	00000000	I:0
тср	d80b00f0	00180012	0a120383	I:0					
RST	40000000								
SYS	00000000	000000e8	000000000	00000000	000000000				
No messages									

Figure 4.7: A console-based view of IP-core registers and interrupts used in development systems is available for monitoring and controlling the ONSEN system. With the introduction of EPICS the application remains for debugging purposes and handles certain initialization procedures.

is found out and the configuration is adjusted before the network is started. This feature makes the flash configuration obsolete.

The reliability of flash content is limited if Linux is not shut down gracefully. The reprogram capability of the firmware by the CPLD introduces the risk of a broken flash chip resulting in lost configuration and license files for SiTCP. With changes in IPMI logic, the firmware reprogram procedure is preceded by a command from IPMC or MMC on the serial console. This command shuts down Linux. This graceful reprogram can also be issued via slow control.

4.3.4 Status

The full ONSEN setup has been successfully installed in the electronics hut adjacent to the Belle II detector. This setup comprises various components, including a debugging PC, optical cables, and the ATCA shelf. The debugging PC plays a crucial role in troubleshooting during operation and is used for firmware and software updates. It is connected via JTAG and USB to all FPGAs. To ensure seamless data transfer, optical cables have been connected to DHH and EB2. Additionally, the ATCA shelf, which houses the ONSEN system, is powered by three power supplies. This redundancy ensures an uninterrupted power supply and minimizes downtime. Two shelf managers work in parallel and remotely control and monitor the ATCA shelf. All xFP cards and CNCBs are attached via Ethernet to the slow control network.

As Phase 3 of the Belle II experiment is limited to the reduced PXD, only 5 CNCBs and 17 xFP cards are powered and used in production. However, the remaining



Figure 4.8: Photograph of ONSEN system installed at KEK. Optical cables in light blue are routed to patch panels that connect to EB2 and DHH. Ethernet cables in yellow and red connect to a switch above for slow control.

boards are kept updated with the latest firmware and serve as spares in case of failure. When the full PXD is installed, additional spares will be shipped to KEK to ensure reliability of ONSEN.

4.3.5 Future Development

During Phase 3 ONSEN proved to handle current data rates easily and helped to identify issues with DHH and HLT. Improvements in firmware and software supported online monitoring significantly for PXD shifters and accelerator experts. The current system runs stable and no changes are expected. However the number of spare hardware is limited and no additional CNCBs or xFP cards can be produced. An updated version of the CNCB has already been developed by IHEP. This might serve as a replacement with only a few changes necessary to the existing xFP cards. However the CNCB firmware cannot be adopted without major changes.

Additional features have already been considered to be implemented at firmware level. However, due to the SiTCP core on xFP cards, FPGA resources are nearly exhausted and changes must be applied carefully to maintain stability.

Chapter 5

Slow Control

5.1 The EPICS Ecosystem

The *Experimental Physics and Industrial Control System* (EPICS) is a compendium of several software tools that allows scalable connectivity between multiple independent hardware and provides distributed control mechanisms to operate accelerator or detector machinery. Each hardware component requires network connectivity, which can span different network layers. This section focuses on the infrastructure of EPICS related software, different mechanisms in use and its control and monitoring possibilities.

All software uses the same protocol to exchange information, called *Channel Access* (CA) [63, 64]. Typically a single Input/Output Controller (IOC) handles single hardware or is dedicated to a certain purpose. With its own database an individual set of process variables (PVs) can be created and altered while running. Such PVs can either represent a specific value contained in hardware, a prepared configuration or a calculation of other PVs. A client such as another IOC, a command line interface (CLI) or graphical user interface (GUI) software can connect to the IOC server, retrieve information or subscribe to future updates. EPICS IOCs publish their information automatically to all registered clients on change. However this behavior can be adjusted by the definition of each PV.

Control System Studio (CSS) [65] is an Eclipse-based GUI window design tool for presenting PV values as text or graphs over time. Individual *Operator Interfaces* (OPIs) can be developed and displayed within CSS. While PVs only provide real-time information previous data must be stored on an additional archiver application [66]. With CSS's *Data Browser* plugin, trends can be displayed using archiver data. Other plugins worth mentioning are the alarm system including audio output and access to a configuration database that provides additional PVs.

EPICS is used for all control and feedback interaction for the PXD, which runs in parallel with data acquisition. Each ONSEN hardware runs an IOC on its PowerPC and servers provide the necessary infrastructure. To provide an easily accessible, fast, and intuitive way to access the system, extensive progress was made during



Figure 5.1: Main operator interface for the PXD, providing helpful information like the average occupancy seen in each module, the status of the power supplies and the data acquisition system. Buttons allow to navigate to further views to get more detailed information for shifter and experts.

the work on this thesis including the unification of OPIs. Adjustments that have been crucial in stabilizing the system, reducing failures and downtimes as well as ensuring safe operation will be highlighted in the next sections. Some features come with a CSS illustration. In figure 5.1 the main shifter interface is shown.

Logging

The slow control also includes a logging system running separately from EPICS. Most IOCs are able to send log messages to the server to track events or inform the shifter about significant changes or problems. It contains a descriptive message and a time stamp, and is assigned a severity or importance. All log messages are stored in a database and can be displayed in CSS. Each IOC can choose a minimum severity for each output, e.g. a lower one to write debug information to a file and a higher one to be sent to the database. Various severity options are listed below from highest to lowest.

Name	Use case					
Severe	possible hardware damage, requires immediate action					
Warning	unexpected behavior, requires attention					
Info	inform about expected state change					
Config	configuration changed (automatically)					
Fine	default internal changes, start of sequences, etc.					
Finer	1st debug level					
Finest	2nd debug level					

5.2 ONSEN Slow Control

Slow control is fully implemented with EPICS on the ONSEN system. Nearly all IP cores within the firmware provide slave registers that are mapped into the filesystem of the operating system running on the PowerPC. The PlanAhead software by Xilinx [67] allows the developer to create a *device tree*, which is stored in the ELF file¹. An additional EPICS module called *Device Support for Accessing Memory-Mapped Registers* (DevBusMapped) serves as a driver in order to ease access to these memory-mapped registers.

Additional EPICS IOCs run on a server next to the ONSEN shelf gathering information from ONSEN hardware and the shelf itself. While the shelf is controlled by a shelf manager, an EPICS IOC communicates via the shelf manager to the IPMI controller of each ONSEN xFP card and CNCB. This is done in order to read out sensor information like voltage, current and temperature [68] developed for the ONSEN system by B. Spruck. In figure 5.2 an example of the different measured voltages of an xFP card is shown.



Figure 5.2: Sensor information of an xFP card extracted via an IPMI-EPICS-communication and displayed in CSS (developed by B. Spruck).

For debugging purposes all incoming and outgoing data for each CNCB and xFP card passes through the 11_monitor core. Its registers provide two counters, the number of 32-bit words and the number of frames. Such registers are read by the PowerPC at a frequency of 2 Hz. Secondary PVs calculate the data and frame rate. The firmware register provides always 32-bits, but the EPICS library does not support unsigned values. Therefore the highest bit is cropped and the secondary PVs are adjusted accordingly in order to detect overflow and yield useful information.

¹ The *executable and linked format* [61] file includes a header and binary executable data. The processing unit (in this case the PowerPC) starts to read from the given address, where the ELF file is available and executes its code resulting in the start of the Linux kernel with all necessary information and data.

External connections to DHH and internal connections between the CNCB and xFP card are based on the AURORA protocol [69], which provides status information about the link. This information helps determine data flow issues like back pressure. During de-initialization, all connections must be dropped and reestablished when preparing for data taking. The initialization process can only be successfully completed when all necessary links are established. The *Run Control* confirms this by reaching the READY state (see section 5.3).

Many IP cores provide interrupt signals to the PowerPC next to their status registers. While the DevBusMapped driver does not support this, the former *node* program is used to receive and acknowledge the interrupts, sum up and push the updated values to PVs. While in some cases the interrupt is signaled together with more detailed information, the *node* program also reads out a FIFO. This is where error information is mapped to single bits of a 32-bit register. Afterwards such information is sent to the logging system.

An additional feature of the firmware build process is to store the hash² generated by the version control system. This hash is later accessible at a fixed address from the PowerPC. In the ELF file, the hash is stored in a similar way. Both hash values are exported to PVs. Especially during intensive development, comparing both values can prevent incompatibility problems.

Additional data analysis in the firmware (see section 4.3.2) extracts valuable hit information from the data stream. This information is provided as PVs. While each xFP card of ONSEN counts the number of hits since initialization and stores the maximum number of hits since last reset, at EPICS level this information has to be further processed. Each xFP card retrieves only every fourth event from the DHC, but data from up to 5 PXD modules. The information must be gathered and the mean value for each PXD module has to be calculated. Due to independent processing times on each system, timing has to be observed precisely to avoid possible glitches, which can result in spikes or drops. In CSS this is shown in figure 5.3, where for each module the four single values are displayed as well. The extraction of the trigger number is done in parallel. This is where the trigger number of the incoming pixel data is subtracted from the trigger number of the received HLT frame. This trigger *lead* is combined and a rough estimation of the minimum time interval between the arrival of pixel data and HLT information is possible³. During a run it is enforced that the trigger lead is always positive. Otherwise the pixel data arrives after HLT information and the pixel data will not be processed anymore. If the time interval is too long, the memory of the xFP card may reach its limit of 1.5 GB and back pressure will slow down the overall trigger rate of the whole Belle II detector. The maximum processing time of HLT is designed to be less than 5 s.

²The hash is a calculated value of fixed length from arbitrary input data. In version control systems like git each development step can be committed and is assigned a unique hash value. In most cases the first 8 characters of the hash are already sufficient to identify the version in use.

³The arrival of the HLT information includes the processing of detector data, creation of RoIs and transmission to ONSEN (including the internal transmission from the Merger node to the Selector node). The arrival of the pixel data depends on the event building on DHE and DHC level including possible timeouts, where the DHP is stuck.



Figure 5.3: Calculation of average pixel per event per module over multiple Selector nodes. The average occupancy results from dividing by the maximum number of pixels per module ($192\,000$). The plots are divided into forward and backward modules. The fluctuations for all modules are caused by the injection pattern of the accelerator.



Figure 5.4: The trigger lead is the difference between the trigger number of the pixel data and the trigger number of the ROI data. Since all Selector nodes run independently the EPICS PVs are not published synchronously and differences are possible. Therefore the maximum, minimum and mean of the differences are calculated. The minimum value is divided by the trigger rate to estimate the minimum time interval.

5.3 ONSEN Run Control

The run control (RC) runs next to the slow control. It's used to steer individual systems and manage measurements of the whole Belle II detector, in order to react on the status of those systems accordingly. In order to achieve unification, all systems must follow a precise scheme, in which distinct states can be taken according to particular circumstances. Such conditions are defined for each individual system, and only a few general criteria must be met. In figure 5.5 a state diagram is shown, including all possible states of ONSEN and its subsystems. The state diagram is the same for PXD and other subdetectors of Belle II. For debugging purposes it is possible to ignore the requested state from Belle II and follow an internal request state.



Figure 5.5: State diagram of the ONSEN run control with stable and transition states. The current state is based on all subsystems, each possible to define individual criteria, and the requested state from Belle II. In this example the current state is RUNNING indicating active data acquisition. Colors are used for each state for better visualization for the shifter. On the right all subsystems (CNCB and xFP) are shown with their own RC state.

There are three stable states: *NOTREADY*, *READY* and *RUNNING*. While in NOTREADY all connections are dropped and data transfer is inhibited, the transition to READY includes initialization of the firmware and the establishment of internal and external connections. Only when all connections are up the READY state can be reached. Otherwise the responsible subsystem remains in LOADING. If any connection drops after READY has been reached, the state changes to ERROR forcing the entire run control of ONSEN and PXD to ERROR as well. In this case the initialization must be issued again. Further configuration of firmware registers is also executed during LOADING. Those registers are read back continuously once per second and if the desired value is not obtained within a timeout of 5 seconds another ERROR condition is met. The initialization also includes the reset of all firmware cores, clearing of the LUT and the memory map and storing a DHC dummy

event in the memory at a fixed address. This dummy event is sent out in case HLT information for a given trigger number reaches the Selector node, but no pixel data was received earlier. The same mechanism is in place for the Merger node, when a dummy event of DATCON is used instead. Most internal connections between cores are equipped with a logical gate preventing data flow while not in READY. All gates except for the *Belle II Format Handler* core are opened. This way ONSEN can receive pixel data without running into full memory when performing debug runs without HLT input.

When RUNNING is requested, internal counting mechanisms are reset and incoming data is passed on for further processing. This helps to prevent unexpected incoming data to get stuck somewhere within the firmware. More complex criteria must be met for the transition from RUNNING to READY. During STOPPING the data rate from the DHH is expected to reach zero and the memory occupancy has to fall below 10 %. The latter constraint is used to make sure that there is enough memory available for buffering when data acquisition is restarted. Otherwise an ERROR state is enforced to clear the memory completely at next LOADING transition. The READY state is reached again as soon as no incoming data is detected for at least 2 s. At the same time each subsystem logs a summary of detected format errors of the incoming data over the whole period in RUNNING. In addition, the number of format errors per second is calculated. Reaching more than 100 errors per second causes an ERROR state in order to prevent further data acquisition.

5.4 PXD Slow Control

Besides ONSEN further improvements have been made for other systems of PXD to handle alternating behavior of the PXD modules, stabilization of existing mechanisms and handle hardware failures in a more automated and faster procedure leading to drastically decreased downtimes of the PXD detector. At first, innovative mechanisms will be presented followed by improvements to existing software.

5.4.1 Overview

The PXD slow control is a set of IOCs, each serving a dedicated purpose trying to disentangle each part as efficiently as possible. Where this is not feasible, handshakes are used to link individual IOCs. The following schematic (figure 5.6) illustrates the main IOCs and their interconnections to provide a first overview.

During operation additional IOCs are running, monitoring and controlling the main power supplies, the crate controller and fans, the optical high-speed link switch and for steering the local DAQ server and its data taking. The DHC, DHI, DHE and PS IOC are common IOCs not part of the power (HV) system or the run control (RC). Other IOCs in the schematic are either part of HV or RC and operate independently. Both systems follow a tree structure, where each IOC has its own state and upper level IOCs combine the states of their children. Further details on the PXD RC part will be given in section 5.4.10.



Figure 5.6: Overview of the PXD IOCs and their connectivity. There any many smaller IOCs available and the counter next to an IOC does not necessarily represent the number of actual applications started on a server. This schematic focuses on the IOCs with highest priority and are essential for the minimal operation.

The PXD module is powered by the power supply, which is controlled by the *PS IOC*. Its communication with the DHE and the ASICs is handled separately by the *DHE IOC*. Since voltages and currents must be applied in a safe order, a state machine, called *PS Sequence*, takes care of extensive checks on applied voltages and makes sure that currents operate within their limits. During powering up of the module several configuration steps of the ASICs are necessary, which are handled by an additional state machine, the *DHH Sequence*. With the DHH Sequence and PS Sequence running independently, a managing state machine, the *PS Control* IOC, calculates a combined state from the DHH and PS Sequence and steers both sequences simultaneously. The default operation mode is to request a new state by the shifter. PS Control forwards this request to all DHH and PS Sequences that are selected for inclusion. When all modules have reached the requested state the current state of PS Control also updates and the transition is complete.

Similar to the run control, there are three stable HV states: *OFF*, *STANDBY* and *PEAK*. Transition to a "lower" state is always possible for safety reasons. In figure 5.7 all main states are illustrated. Detailed steps on the ASIC configuration is given in section 5.4.9



Figure 5.7: State diagram of the PS control with possible request states in gray and additional transition states. On the left side the meaning of the corresponding stable state is given.

5.4.2 DHP Link Recovery

The PXD modules send the zero-suppressed data via high-speed (HS) links to the DHE. Each DHP has its own connection and the data is merged by the DHE firmware. Receiving data via these optical HS links is impaired by different sources. The transceiver receives the optical signal and converts it to an electrical signal using multiple amplification mechanisms. These mechanisms may not be fully optimized and the resulting signal cannot always be correctly translated into a binary sequence. At a certain point, the link quality is so severely limited that this link is considered broken. To recover the link, re-initialization via DHP logic is unavoidable. The DHP sends a key word to synchronize the link after reading out one frame, which corresponds to the entire matrix. A bug in the logic of the DHP can lead to the fact that with large amounts of data this key word no longer arrives at the DHE within a specified time. In this case the link is considered dead, so re-initialization of the DHP is necessary.

During the operation of the system, it has also been noticed that under certain conditions the DHE no longer accepts data from the HS Links. While the link is still up and alive the DHE uses empty frames to merge its events. The so call *ghost frames* increase in number, and the DHP reset helps to recover from them. It has been developed an IOC to monitor all links to ensure that they are immediately recovered in the event of any issues described previously. For events with high occupancy, the rate of links being restored increases. This way it is not necessary to start a new run and data acquisition is only marginally affected. On average, a link is restored within 0.2 s.

5.4.3 Module Recovery

The power supplies of the PXD modules have an over voltage protection (OVP) extension board attached since the start of Phase 3. It is designed to monitor a subset of all 24 available channels and verifies that the voltage at load operates within certain boundaries. In 2022 it was discovered that its logic is not radiation hard and the board can trigger an interlock signal to the main logic of the power supply. The module is shut down immediately, all voltages are not applied anymore and no regulation of currents is performed. As a consequence the readout ASICs and the matrix are disabled and data taking must be suspended.

The module being shut down has to be powered again and the configuration of the readout ASICs re-initialized. These processes are executed by the *DHH Sequence* and *PS Sequence*, two IOCs that have all necessary steps defined in their state machine. However the sequences have to be started manually by the active shifter, who is informed by a visual and audible alarm.

Here a new IOC jumps in and recovers the module back to its previous state automatically. The OVP interlock signal is cleared and the module is powered again with full configuration and previously used pedestals within less than 2 minutes. The detection runs fully automated by checking the combined state of all modules being in TRIP⁴. With this mechanism PXD is the first subdetector of Belle II to support automatic recovering without any shifter action required. Special care must be taken in case the sequences do not perform their desired work due to internal error conditions, e.g. a voltage is not reached anymore or the configuration fails. Especially in the beginning of this development the pedestal upload did not work reliable and further checks had to be made to ensure the module to be working properly again. In addition, time could be saved during ASIC configuration by postponing the pedestal upload so that it is timed with the power-up of the matrix. Since powering up the matrix takes longer than uploading the pedestal, about 50 s are saved for the entire module recovery.



Figure 5.8: Simplified state diagram of the module recovery. The same state machine is running for each module individually.

⁴The TRIP state by PXD is used differently compared to other subdetectors. While the state is intended to indicate an overvoltage, the state is used by PXD in case fewer than 3 modules have been shut down by any interlock signal at the same time. Details can be found in subsection *PS Control Update* below.
5.4.4 Test Pattern Upload

While the accelerator is being optimized and not available for data taking, test runs are performed by Belle II with the designed trigger rate of 30 kHz. These are intended to help identify problems in the overall structure of the data acquisition and its associated systems improving its stability in the long term. During this phase, most subdetectors, like the PXD, are not fully operational. The sensitive matrix of the PXD modules is not supplied with voltage and the readout electronics cannot produce any meaningful information. Since noise on the electrical lines can lead to very high and unpredictable data rates, it is possible to upload a predefined pattern into the ASICs' memory. For this purpose, the readout of the DCDs can be stopped and the memory content of the DHPs remain constant.

During Phase 2 and in the beginning of Phase 3 it was recommended to join specific tests runs and uploading a test pattern required manual interaction by an experienced shifter. A Python script was supposed to be executed on the local DAQ server while the modules are supplied with voltages for the digital part. When the accelerator prepares for physics runs and the matrices of the PXD modules are powered the configuration had to be reverted manually again.

In order to enable this feature for all PXD shifters a new IOC was developed to create and upload a test pattern for all modules with a given occupancy level to simulate higher or lower data rates. Higher data rates help to provoke back pressure throughout the full chain of DAQ systems and allows to verify that all individual systems behave as expected. In case the modules are ramped up again, the test mode is automatically reverted and most recent pedestal are uploaded again. This is necessary, because the pedestals have to be invalidated during the operation with a test pattern in order to not have masked pixels left over that do not contribute to the desired occupancy.

5.4.5 DHP Register Monitor

Due to the short distance to the interaction point the PXD detector and all its digital components are exposed to a heavy amount of radiation. The internal registers in the ASICs are hardened against single event upsets (SEU). Unfortunately it was discovered by chance that the implemented mechanisms do not work reliable under the circumstances of increased background rates.

To evaluate the gravity of the situation repeated checks of the registers are necessary. A new IOC checks all registers periodically and compares those against their previous value. The registers of the DCD are skipped during this part, because they provide only a writing mechanism via their JTAG connection. The switchers require the left- or rightmost DCD to be included in the JTAG chain, which can induce additional noise during active data taking and are therefore skipped too. Finally 1756 channels of all DHPs are read and compared every 20 s. Log messages are submitted when changes are detected.

The standard shifter is not capable of recovering such SEU instead an expert shifter is consulted. Therefore at a later time the IOC was reworked giving it the

ability to revert the register value and write the corresponding register via JTAG fixing the configuration of the DHP. Still the shifter is informed about an SEU, but not further action is needed and the log message appears in the shift report at the end of the shift as well. As part of this rework the DCD registers are rewritten as well increasing the total number of registers to 3024. During every cycle only a single DCD configuration is applied to reduce the impact of the changed JTAG chain. Writing DCD registers automatically alters the JTAG chain, afterwards the JTAG chain needs to be reset only including the DHPs.

SEUs due to radiation should only occur at a very limited rate. In case of multiple register changes the IOC shuts down the module completely and makes use of the *module recovery* to power it up again.

5.4.6 Clear Channel Handling

After a beam incident in November 2019 one module (1.8.2) changed its behavior visible via the current drawn by the *clear* channels. After the next power cycle the module could end up in one or another state. Uploading pedestal data taken not in the same state can result in high noise. It turned out that it is possible to force the module into one of the two states by adjusting the *last_row* parameter of the DHP configuration back and forth. While the value is set to 191 by default, a lower value causes the switcher to skip the last instructions of the switcher sequence, which is stored in the DHP memory. The "ser-in" signal at the very end is skipped and the switcher does not continue to switch gates. This leads to low voltages of the *gate* and *clear* channels. A mechanism was added to always ensure the correct *clear* current and its state after the module has been fully powered. Later in 2020 the module did recover from this behavior and the mechanism was deactivated.

5.4.7 PS Control Update

The main steering application for all PXD modules is the *PS Control* IOC. It interacts with the global high voltage control (HVC) system by Belle II and the two PXD sequences for the DHH system and the power supplies. By requesting a new state, either by the Belle II HVC or manually, the transition task is forwarded to the sequences and all selected modules are powered step by step individually. When all selected modules reach the requested state the current state of PS Control also makes the transition to the requested state.

Including or excluding a single module is possible only if the current state of PS Control and the module itself match. For each state of PS Control only certain states of the modules are possible, because not every transition is allowed. As soon as any module steps out of line, the current state of PS Control is returned as UNKNOWN. A special handling is in place in case any module reaches the ERROR state, which has priority and the state of PS Control switches to ERROR as well.

The IOC was tweaked during Phase 3 to react more quickly in certain situations and to converge with the Belle II scheme. With the installation of the overvoltage board, individual modules were abruptly deactivated and had to be powered up and configured again. However, the PS Control does not easily allow the state change of a single module, so consequently all other modules had to be deselected and the one module could be restored to its initial state with the PS Control. With the help of the request state for all modules and a newly added request state for each individual module, the recovery could be accelerated. After deselecting the problematic module, its request state could be adjusted similarly to the overall request state. Later, this execution was adopted by the recovery IOC and this functionality remains of interest only to experts. In 2020 the functionality of powering up a single module was extended. Since the overall state of the PXD is visible to the Belle II shifter, no module may enter a higher state at any time for security reasons.

The Belle II high voltage scheme includes a TRIP state, which was unused for PXD, because no behavior of a tripping voltage is possible. Instead, the state is used for module recovery. In case of failure of up to two modules, the overall state of PXD changes to TRIP. This implies an automatic recovery to the previous state without action of a shifter.

5.4.8 PS Sequence Update

The *PS sequence* controls all channels of the power supply of each module by following a fixed order. Current limits and voltages are set one after the other so that the PXD module never suffers any damage. In coordination with the *DHH Sequence* the configuration of the ASICs is applied at the same time during powering. The execution order was not changed in the context of this work, but individual steps were extended in order to extend the safety of the modules.

An important step from first characterization in the laboratories to the productive use of the modules as part of the Belle II detector is the verification of the set voltages. This includes the state of the current regulator so that no channel is operating in current limit. After almost every single step of the sequence, the previously set channel and related channels are verified. The measured voltages are calculated using an ADC inside the PS and some tolerance must be allowed. Since each ADC is built for a different voltage range, the tolerance is also adapted to it. With three times the step size of the voltage value, the voltage can be checked sufficiently precisely.

Additionally, these checks were made more robust. It was ensured that the transition from a higher state to a lower one is always possible instead of blocking any further operation during powering. Before the change an emergency shutdown could not be caught by the sequence in some cases under certain conditions, which could lead to an unavoidable restart of the software in the worst case. A similar problematic behavior was prohibited where the micro controller of a power supply was suddenly restarted by an internal error after or during startup, causing the module to shut down.

In the course of the increased radiation dose, some of the previously described checks had to be adjusted. There are also differences whether a module is powered up or down, as leakage currents may flow through different pathways. One of the most significant problems of radiation is the highly increased current flow in the hv channel. Previously performed irradiation campaigns had not shown a current higher than 1.3 mA. Before the PS were modified several times during Phase 3 to allow a 20 times higher current, the PS sequence had to be adjusted so that the check of this channel would not block powering.

Another special issue of the power supply is an error of the overvoltage protection (OVP) board. If the time between shutdown and powering of the modules is too short, information on the I2C bus can not be read correctly by the micro controller and thus the power supply prevents from being turned on. A delay as well as an additional sequence for the complete restart of the power supply via the crate controller provides increased stability at this point.

A similar feature required manual intervention. After switching on the power supply, there were occasional problems with the hv channel, where the maximum possible voltage of -80 V was measured, although no voltage had been set yet. To prevent damage to the matrix of the module a check was added at the beginning of the sequence to switch the power supply off and on. This procedure does not result in a state change in order to handle it quietly without being noticed by the Belle II shifter. Only the PXD shifter is informed by a log message, which also allows experts to keep track of this issue.

5.4.9 DHH Sequence Update

The DHH sequence works in parallel to the PS sequence and is responsible for DHE, DHI and DHC. At the beginning the configuration is applied from the database to DHE, DHI and DHC. This ensures that possible previous changes by local work are reverted. During power-up of the modules the configuration of the ASICs is set, because without applied voltage the ASICs lose their configuration parameters as well as the content in the memory. After writing all registers of DHP, DCD and switcher, a special sequence for the switcher is uploaded into the memory and verified. In combination with the sequence of the DHI, the high-speed connections between DHP and DHE are reset and checked. In figure 5.9 a state diagram is shown including the main steps.

At the beginning of Phase 3, the DHH sequence suffered from massive performance drops each time at the start of powering up. After an extensive review of the source code, large parts of the sequence were rearranged. The sequence is built in a modular way that all necessary PVs are connected based on the module name. While EPICS is optimized to ensure permanent and reliable connections, breaking and reestablishing such a connection requires more resources. In the source code, all connections were reestablished every time the modules were powered, while one connection setup after the start of the IOC is sufficient, since the PV names do not change later on. To ensure flawless execution of the sequence during powering, all PV connections are verified at start.

The evolution of the DHE software enabled new and improved strategies during the application of the configuration. For example, a PV was introduced that confirms



Figure 5.9: State diagram of the DHH Sequence. The rectangles represent persistent states, the diamond shapes transition states. The double arrow transitions require a new requested state. The inverse way is also possible, but nearly without any action. A transition from every state to an ERROR state is possible in case the desired action failed or the connection to the IOC or hardware failed.

the last access to the memory of the ASICs as successful. The DHH sequence was adapted so that this PV is checked each time data is uploaded to memory. This increases both the speed of the upload and the security of correct data in the memory. Especially the switcher sequence must be uploaded correctly. Otherwise gating and clearing operations of the switcher may harm the matrix or the ADCs. Another feature of the DHE software allowed a bulk read of the ASIC memory, which is much faster instead of reading single 128-bit words. This mechanism required additional changes in the DHH Sequence and were also applied to the memory verification of the offset data and pedestal data.

The galvanic isolation between DHE and DHP was successfully introduced with the DHI after Phase 2. However, it is necessary that this connection is also interrupted after the shutdown of the modules. It does not matter whether the module is powered down normally via the sequence or the module was quickly shut down by an interlock signal or emergency call. A similar approach was taken in measuring the temperature of the DHPs. Once the configuration of the ASICs has been completed, the temperature must be checked regularly to prevent overheating. While in Phase 2 this was still done manually for 4 modules by the shifter, this mechanism was automated for Phase 3. Later in the development of the DHI software, DHP temperature measurement was made possible in parallel with the configuration of the ASICs⁵.

To ensure the stability of all systems during data acquisition, dummy triggers are used to read out all subdetectors during periods without colliding beams or during beam studies. This simulates test conditions at maximum trigger rate so that potential problem spots can be eliminated in advance. For PXD in this case, a test pattern is loaded into the memory of the DHPs using the IOC described in section 5.4.4. During this time, the matrix of the PXD modules is inactive to prevent possible damage from radiation effects. The DHH sequence has been extended to automatically detect if the DHPs are in test mode and to automatically reverse it in case the matrix is being activated. Due to the reduced bandwidth⁶ between DHE and DHC only half of the design occupancy could be tested at 30 kHz.

Similar to the use of the test pattern, an alternative was previously developed to participate in test runs. The size of the memory in the DHP allows to store 2 different pedestals. While an 8 bit value is stored for each pixel, the maximum value of 255 is assigned a special function where the corresponding pixel is completely masked. If all pedestal values are set to 255, the whole matrix is masked and no hits are sent from DHP to DHE. This allows further event building in the DHE at 0 % occupancy, allowing higher trigger rates up to 50 kHz. Later in Phase 3 it was decided to not participate in those test runs after it was proven not being the bottle neck in any previous tests.

While runs with test pattern or masked matrix have lower priority, data taking under real conditions is much more crucial. The pedestals are generated separately with the local data acquisition (DAQ). Everything is controlled by the Calibration IOC, a software part of the Calibration Framework (see section 5.5). The Calibration IOC also handles the upload of the pedestal to the ASIC memory. In order for the configuration of the ASICs and the upload of the pedestal via the Calibration IOC not to lead to any complications, both IOCs must coordinate themselves optimally. Especially during the power up of the modules, these processes were optimized several times during Phase 3. Finally, the upload of the pedestal was shifted to the phase in which the ASICs are already completely configured and the matrix voltages are applied. In this way, the duration of the entire start-up process could be greatly reduced in order to get the detector ready for operation as quickly as possible.

⁵The temperature measurement of the DHPs is done by setting registers in a specific order. Since all register accesses are transmitted from the DHI to the DHP via JTAG commands, measuring the temperature and uploading the configuration at the same time can lead to unpredictable conflicts. With new DHI software in the course of Phase 3, this issue was overcome with improved prioritization and the temperature measurement is turned on as soon as the DHPs are supplied with their operating voltage.

 $^{^6 {\}rm The}$ board design of the DHC suffers from low signal amplitude on the transfer lines between the DHC and DHE. The amplitude of the eye diagram is not sufficient for 5 Gbps transfer rate when connecting 5 DHEs simultaneously. With reduced rate of 2.5 Gbps the DHC works stably. A new version of the DHC is expected to be ready in 2023.

5.4.10 PXD Run Control Update

The *PXD Run Control* describes a part of the PXD slow control that deals exclusively with the coordination of runs. While separate IOCs are used for the DHH, ONSEN and DATCON subsystems, they are all combined and controlled by a higher-level IOC. All follow the same finite state machine, which is illustrated in figure 5.10. Each state machine is also implemented in lower levels allowing individual boards to be masked if necessary.



Figure 5.10: State diagram of the run control with possible request states in gray and additional transition states. The ERROR state can be entered for all states.

A separate system that does not follow the transitions of the finite state machine is specifically incorporated. The Calibration IOC as part of the Calibration Framework (see details in section 5.5) establishes a connection with the higher-level run control via 3 PVs. The following 3 transitions are blocked as long as the corresponding PV is set to 1.

PreventLoading / PreventStarting

prevent transition from NOTREADY to LOADING or from READY to STARTING

The Calibration IOC allows the run control to disable this transition, because during LOADING and STARTING the run control of the DHH system applies configuration to the system. The Calibration IOC instead changes the configuration of the DHH system in order to take new calibration data. This handshake prevents interference and the higher-level run control has an internal timeout of 20 s before an ERROR state is returned.

PreventRunning

prevent transition from STARTING to RUNNING

The Calibration IOC takes care of uploading pedestal values into the ASIC

memory. With the transition to RUNNING the trigger system of Belle II is not on hold anymore and data acquisition could be started with invalid calibration data. The run control waits up to 90 s for the upload to complete before returning to ERROR. The high timeout is necessary in case the upload failed once and the Calibration IOC tries again.

5.5 PXD Calibration Framework

The second part of PXD software related to module operation and characterization is the *Calibration Framework*. Started by single files for calibrations by the laboratories of PXD it evolved to a comprehensive collection of scans and IOCs to ensure best performance of PXD modules before installation and during operation. The framework is written in Python and uses the *PyEpics* module [70] to interact with the hardware as well as the *PCASpy* module [71] to create EPICS like IOCs. During the work of this thesis different aspects of this software have been adjusted to improve massively stability and handling of most parts of this framework.

The framework is developed by the same version control system (Git) used by all other software parts for PXD. While the necessary IT infrastructure is hosted at DESY, the offered tools also provide applications for continuous integration and automated builds. With the framework mainly written in Python, its functionality can be tested and verified using the *pytest* module [72]. With over 700 tests of scans, analyses and other functions with varying parameters the framework's stability has been greatly improved, which also aids for future changes. After the transition from Atlassian tools [73] to GitLab [74] the test results are automatically published allowing all developers to directly see new problems. An example is given in figure 5.11. All tests are executed for each change in the default branch and for each merge request⁷ ensuring that future changes will most likely not break existing functionality and not introduce incompatible source code. The tests for scans have been extended to work with the former Hybrid5 modules as well, which provides a reduced matrix and smaller events read out from DHH [75, section 4.8]. With a reduced delay after each EPICS communication and adjusted scan parameters the tests of calibrations only require a fraction of the typical duration. With all tests split into categories, multiple tests are executed simultaneously and the average testing duration could be reduced to 10 minutes. The tests are executed within a Docker [76] container, which is created once with all dependencies. Since reproducibility plays a major role, all necessary files are part of the repository.

Further improvement of handling modules was achieved by the introduction of an abstraction layer. Originally introduced by H. Schreeck this layer allows to change parameters of one or multiple modules at the same time. Several enhancements have been applied during this work. With delays after specific actions and optimized access to the EPICS layer, glitches during applying the configuration could be ruled out

⁷A merge request can be used by a developer, when committing changes to a separate branch. The request aims to merge the new or changed code into the default development and allows developers and maintainers to review and discuss its content.

5.5. PXD CALIBRATION FRAMEWORK

Summary 1466 tests	0 failures	0 errors	10	00% success rate		2340.95s
Jobs						
Job	Duration	Failed	Errors	Skipped	Passed	Total
test-3-cal-h5: [el7]	269.41s	0	0	1	22	23
test-2-cal: [el9]	478.89s	0	0	2	38	40
test-3-cal-h5: [el9]	272.87s	0	0	1	22	23
test-1-lib: [el9]	394.27s	0	0	2	668	670
test-2-cal: [el7]	517.73s	0	0	2	38	40
test-1-lib: [el7]	407.77s	0	0	2	668	670

Figure 5.11: Test summary of a GitLab pipeline. The pipeline (execution of multiple jobs for each commit) is divided in 3 jobs that run in parallel. Each of these jobs is executed for two environments that use different version of OS, Python and Python modules. This prevents unexpected behavior for future upgrades.

completely. Cross checks for ASIC configuration and checks for voltages of the power supplies have been implemented increasing the safety of modules. Furthermore all voltage changes are converted into ramping procedures with safe step sizes. While ramping, the current consumption of the selected channel and the *source* channel are closely monitored in order to prevent operation in constant current mode.

With the abstraction layer further improvements were possible for each scan. Scans are divided into three parts. In order to unify individual calibrations and create new ones in a more structured way with as little duplicate code as necessary, specially designed classes were introduced. These classes were used to do most of the preparation. This included the unification of the folder structure as well as the configuration of the readout system. Further the analysis of the data was parallelized in such a way that the number of logical processors is utilized as efficiently as possible by employing a pool of workers. At the time of writing 14 existing calibrations have been ported to the upgraded scheme and make use of the abstraction layer in any conceivable way. Using the testing procedure mentioned earlier, all of these calibrations support code testing, most of them also for Hybrid5. Before testing, an emulator of a PXD system is launched, which provides all PVs without connected hardware and can simulate most behavior. With its help, the general communication via EPICS as well as the entire execution of a scan can be tested. This also assists in the development of new scans allowing their execution in a safe environment without risking module damage. All scans are divided into the following stages:

Measurement

During a measurement the system is swept over a certain predefined set of configuration parameters and data or currents are recorded. Afterwards the configuration is reverted and the initial state of the system is restored.

Analysis

The analysis can be executed independently of the measurement, where all necessary information has been saved. No interaction with the actual hardware is done at any time. This is especially helpful in case the analysis is more complex and requires more time to complete. Similar to the measurement, the analysis is configurable and results are stored as data files with new configuration parameters or PDF files included images and plots.

Update

The resulting configuration parameters can be uploaded to the active system or to the configuration database.

These classes deal with the problem that multiprocessing in conjunction with EPICS connectivity does not work as well as expected. When using multiprocessing in Python on Linux the existing main process is forked, which results in a new process with duplicated memory. This memory also contains sockets for the TCP connection used by EPICS PVs. After being forked, these connections must be reestablished within the new process, which requires a cleanup at the beginning. The PyEpics module provides a specially designed subclass for this purpose called CAProcess. Due to a misbehavior of the underlying library the EPICS cache containing all connections can not be cleared completely in rare situations. This leads to processes that get stuck and the measurement or analysis fails and has to be repeated. When a fresh calibration is requested between runs and there is only a short amount of time available, this delay is particularly problematic. With the addition of classes and the abstraction layer the measurement stage has been rewritten in order to rule out multiprocessing completely without performance degradation. The same approach cannot be used for the analysis, since the processing speed would be severely affected without multiprocessing. Instead, a pool of worker processes is started at the beginning so that no active EPICS connections need to be closed. These worker processes are idle most of the time and jobs are submitted once an analysis is initiated.

5.5.1 PXD Shifter IOC

In addition to the general aspects of the framework, special software was written in order to automate the duties of people taking shifts for PXD. The goal is to offload as much work as possible from shift workers. This involves automating individual tasks that have to be dealt with at the beginning and end of a shift. A welcome side effect is the reduction of wrong or unspecific information in shift reports due to missing notes or comprehension problems.

At the beginning, the shifter has to announce himself in the chat, so that it is clear for other shifters who can be reached in case of trouble. The data entered for the user name and telephone number are saved locally, so that only the name has to be entered for the next shift. The shift table for the next day is extracted with the help of the shift schedule web page. Therefore, the name of the next shifter and his personal data are automatically prepared when the shift changes.

During the shift, the quality of the measured data is to be checked by means of automatically generated histograms. A part of this evaluation is provided as EPICS

PVs, but this task is not yet completely taken over by the PXD shifter. The latter must manually enter the rating of each run into a database. The entered data is later extracted from this IOC for the shift report that must be prepared at the end of a shift.

The shift report gathers data from multiple sources. The evaluated data quality of the runs, the messages in the log, and multiple PVs are constantly monitored to generate the most detailed report possible. Finally, this is uploaded to the electronic logbook and only minor adjustments are required from the shifter. These tasks take a huge load off the shifter and help to keep the shift reports consistent. With timestamps for every change, it is more convenient for experts to review the shift. In table 5.1 an overview is given of how many shift reports have been created before and after the introduction of this IOC and its capabilities. Manually created reports after November 2020 can have multiple causes. A network connection may prevent gathering all the necessary information or the shifter was not aware of the automated process up until now. While the report is created automatically, it still needs to be initiated by the current shifter.

Reports	until Nov. 2020	
 manually created 		1125
– missing		111
Reports	since Nov. 2020	
 automatically created 		1394
 manually created 		7
– missing		6
Total Created Reports		2526

Table 5.1: Manual vs. automatic shift reports. On November 17th, 2020, automated shift report creation was introduced. The creation of shift reports was originally added to shifter duties on February 2nd, 2019.

Moreover, for bad runs, tickets are created once per week for experts. The tickets are created in a centralized web service called Jira, which is part of Atlassian tools. Run coordination and PXD experts can communicate by commenting on this ticket to understand the reason for the poor run quality and collaborate to prevent it in the future. In order to start a discussion, the tickets are automatically linked to several experts.

5.5.2 Local DAQ Control IOC

In addition to primary data acquisition, the PXD also stores data locally on a server. In this case, the data stream from one of the DHC outputs is duplicated and sent out via the UDP protocol. This interface does not support the maximum DHC data rate, but allows fast insight into the raw data. This data can be analyzed independently of other subdetectors. However, such data does not contribute to physics analyses based on the data set of the entire experiment. The technique was developed primarily for beam times and laboratory experiments, but still has its advantages. The analysis of the local data during data acquisition allows faster conclusions to be drawn about the quality of the data compared to data quality monitoring of the entire Belle II detector.

This local data acquisition is controlled by the *Local DAQ IOC*. It monitors PVs of PXD and Belle II run control and starts and stops data receiving from all DHCs. At the start and end of each run the system state is extracted from a variety of PVs. This information is compiled into an online logbook entry. After each run the IOC communicates with the DHH run control. Additional raw data is acquired with zero-suppression disabled. The data is analyzed and updated pedestal values are prepared, which can be uploaded to the DHP memory. Performing measurement and analysis in advance reduces the time for newly uploaded pedestals. This allows the update to be scheduled between two runs without further delay.

5.5.3 Local Data Quality Monitor

The application of receiving local data allows to connect via a TCP network socket, where each event can directly be picked up. The *Online Monitor IOC* is analyzing the local data and creates a hitmap and multiple histograms. In figure 5.12 the GUI for the shifter is shown.



Figure 5.12: Example of the results of the Online Monitor IOC. The analyzed data is presented as PVs and are visible for the PXD shifter. By default the histograms are reset for each run and the hitmaps are included in the online logbook entry created by the local DAQ control IOC.

To process the data it has to be unpacked and converted. This code is written in Cython [77] for highest performance. Arrays are then created and modified by the NumPy module [78]. Finally the results are prepared to be published as EPICS PVs. Statistically compiled code is already optimized and leaves little room for improvement. Instead, the entire analysis is started simultaneously in several processes, each for a part of the modules. With the help of pipes the updated PV values are sent to the corresponding IOC process. Performance tests depend on the number of modules to be analyzed. With 19 modules, the typical processing rate is 100 events per second. The first optimization, using pipes to transmit PV updates, increased the processing to over 140 events per second. Profiling was used to identify that a lot of time was spent on channel access communication. Parallelization scales almost linearly with the number of processor cores utilized. Minor performance tweaks were applied to data unpacking.

5.5.4 Data Transfer IOC

Another IOC to increase automation takes care of the downstream processing of the local data of the PXD. This takes place in parallel with the data acquisition of all detectors. The previous procedure was based on a time-based scheduled job to start the analysis for all recent runs using an error-prone script. The potential number of parallel analyses poses a risk to the performance of local data acquisition. Furthermore, compression and backup were manually started by an expert without a schedule.

The first file of the data of each run is analyzed and the results are copied to a web server at DESY. To prevent overloading of the server during ongoing data collection, the maximum number of simultaneous analyses is limited. In the next step, the data is compressed before it is forwarded as a backup to KEKCC, a server cluster at KEK for data processing and storage. PVs for each step show the current progress, and the entire process can be deactivated for maintenance purposes. In figure 5.13 the status during experiment number 26 is plotted as a reference. The IOC was introduced in July 2021 without the analysis step. This was followed by several enhancements and bug fixes to cope with network related changes and information systems for experts in case manual interaction is necessary⁸.

⁸The transfer of analysis plots and backup data requires user authentication. In case of connection drops a user has to reestablish the connection manually.



Figure 5.13: Handling of run data for experiment number 26. Compression and backup are started 7 days after the run has been recorded. Drops in the processed run number are caused by sudden restarts of the IOC. The difference between compression and transfer has been pinned down to the slow network connection of the gateway between the internal Belle II network and the KEK-wide network. The long pause of analysis and transfer starting on June 7th is caused by a network issue.

5.5.5 Calibration IOC

Several optimizations are necessary to improve module performance. Besides the matrix voltages, parameters of the ASICs can be adjusted. This involves changing the configuration for all pixels read out from a single DCD or from one of the three gate areas. Different optimization approaches were developed during the module characterization. However, these must also be prepared for later use. Instead of a single module in the lab, multiple DHCs will be used each receiving data from a total of 5 modules. To transfer this task to the PXD shifter, the Calibration IOC was created, which enables all calibrations via an EPICS interface. Figure 5.14 shows the CSS user interface. During normal beam operation, the pedestal calibration is frequently in use. When more time is available, such as on maintenance days, other optimizations are also performed to compensate for radiation damage.

With the help of calibration classes (see beginning of section 5.5), all calibrations behave similarly and follow the same flow chart. Pedestal calibration is performed with separate configuration parameters to reduce overall time. By default, a configuration snapshot is taken at the beginning of each calibration. This is done in a separate thread, saving more than 2 minutes. Another step skipped is the plotting at the end of the analysis. Due to limitations of the underlying Python module matplotlib [79], this cannot be done in parallel. Instead, the entire analysis is restarted in a new process after the pedestal upload is complete. Plotting would normally take up to 30 seconds per module. In figure 5.15 the common steps for each calibration are summarized. Orange items apply to pedestal calibration only.

A major change has been the automation of performing a new calibration after

PXD Calibration Control					
DHC BUS	6Y	DHE BUSY	CALIBRATION Status Click for PedestalMonitor GateOn Status		
	Inner Forward	🥥 🕑 1011 🔤	recalibration necessary for:		
	Inner Backward	🥚 🕑 1021	need to be		
	Outer Porward	🥚 🕑 1031	please wait		
Н30 🦲	Outer Daonnare				
		IO51 IO51			
H40		1061 1061	Manual Upload		
н50 🦲		1071 1071	Run Pedestal 🥥 Last Run 🗸		
н60		IO81	/data/commissioning/runs/all/EXP0026/pedestal/run2091 Data Path		
			Exp # Pun # Date of selection Upload p p g2022 66 23 322-03 Podestate		
Unselect Phase3	Select Phase3		Schedule Pedestal Upload: ON 🧶		
Select IF	Select IB	Calibration Acti	Calibration Server Status		
Select OF	Select OB	Pedestals in	Advanced > Calibration server started.		

Figure 5.14: GUI of the Calibration IOC. Each module can be manually selected and the current status is visible via the left LED. A progress bar and a status message are displayed for additional information. The upload of previously prepared pedestals can be performed on the left. The *Calibrate in PEAK* option allows a shifter to schedule a new calibration.



Figure 5.15: The individual stages for each calibration. Special tasks for pedestal calibration are indicated in orange. While a calibration is running, another one cannot be started for the same modules.

the modules are powered up to PEAK and uploading fresh pedestals after a run. In both scenarios the module state and the run state are closely monitored. Slightly different states of the power supply channels may result from powering up the modules to PEAK. This can affect the pedestal distribution and make re-calibration inevitable. The difference in power consumption between OFF and PEAK changes the temperature of the module by a few degrees Celsius. The temperature will stabilize within the first 30 seconds. As a compromise, calibration is initiated after 20 seconds and a new run start is delayed. Experience has shown that the common mode shifts during beam operation. With previous improvements in pedestal upload speed, updated pedestals are automatically scheduled after 4 hours. At the next run stop, the upload is triggered with the most recent calibration data.

Chapter 6

Operation Results

This chapter focuses on the performance of PXD and its data acquisition system. Phase 3 data taking began in May 2019 and lasted until June 2022. The operation time was used for machine studies and collision data taking, including summer and winter breaks.

Before reporting details on PXD, a brief summary of the accelerator is given to understand and correlate PXD and ONSEN behavior. The SuperKEKB accelerator started at luminosities below its predecessor and beam parameters were continuously optimized. In order to implement the nano-beam scheme at moderate background levels many efforts were made. Figure 6.1 shows the detected luminosity by ECL through Phase 3.



Figure 6.1: Luminosity of each longer run in Phase 3 detected by ECL. In June 2022 the instantaneous luminosity record of 4.71×10^{34} cm⁻² s⁻¹ was reached [23]. Several plots in this chapter are made using the same time period, allowing for a correlation between the different statistics and the luminosity.

6.1 Instrumental Results

While all other subdetectors were fully installed only half of the PXD was mounted next to the beam pipe due to production delays. However, the full inner layer was present allowing high vertex resolution in all radial directions. PXD participated in all global runs to gather collision data and cosmic runs for alignment studies. Cosmic runs are crucial to ensure high vertex resolution and increase precision in physics analysis. The ONSEN system has consistently remained operational, ensuring uninterrupted data acquisition processes. However, it became necessary to replace two xFP cards within the system without impact on data taking during collision runs. In the following table an overview of the operation time is given.

	Number of runs	Run time [h]
total physics cosmic	$11695 \\ 9453 \\ 901$	$9342 \\ 7454 \\ 1172$

Table 6.1: Phase 3 operation statistics. Given are the numbers of runs as well as the durations where PXD, and therefore ONSEN as well, was included.

Further detailed analysis rely on archived live information. EPICS real-time information is permanently archived and allows to recapture system performance and status. In this chapter many results are extracted from this archived data¹. Additional information, e.g. sensor hitmaps, is generated from the local DAQ data. It receives a subset of data from each DHC for online analysis independently of the Belle II DQM.

In the previous beam period in 2018, ONSEN already demonstrated basic functionality at low rates. At that time only 4 PXD modules were installed and connected to 2 DHCs and 2 ONSEN nodes. The interaction of Belle II, PXD, and ONSEN with the final hardware components confirmed the system integrity. In Phase 3, ON-SEN's performance converged to design values and its behavior and stability were monitored under load.

However, due to the accelerator upgrade the design luminosity is not yet reached and the data rates are still moderate. Figure 6.2 shows an overview of total input and output data rates of ONSEN. Without event filtering the ratio between input and output data will be 1. In 2021, event filtering was permanently enabled, preserving only events with interesting physics data found in reconstruction by HLT. Event filtering was estimated to be at least 3, which was used for performance calculations of ONSEN. This factor is already higher due to prescaling in the level-1 trigger logic and higher backgrounds, as plotted below². Before the shutdown in 2022 studies

¹Plots of the full period of Phase 3 include only data points from physics runs longer than 1 hour. This is to ensure sufficient run statistics and reduce the amount of data points.



Figure 6.2: Comparison between input and output data rates of ONSEN averaged for each physics run. The highest input data rate was reached near the end at peak luminosity in 2022 with over 250 MB/s. In the lower part the fraction is shown. During 2021 the *event filtering* was enabled, which results in a permanent increase of the ratio.

were performed with event filtering disabled, visible by the short drop in the ratio.

ROI filtering

In December 2019, 3 physics runs were verified with *ROI filtering*. Independently of event filtering, zero-suppressed data is analyzed and only hits within any ROI for a specific module remain in the data stream. Figure 6.3 plots the input and output rates of one of these runs. Due to an unusually long processing time, the output data rate is delayed. With event and ROI filtering the data reduction ratio varies between 20 and 30. While event filtering will remain enabled, the ROI filtering activation is still to be decided. This depends on link bandwidth and the total amount of data, because all raw data is stored permanently on tape. Additionally the reduction ratio by ROI filtering is calculated by DQM individually for each module³. With higher luminosity and varying backgrounds in the future, it is not possible to predict the future of the reduction for the time being.

²Incoming and outgoing data rates are directly compared in the plotted ratio. Without event filtering the outgoing data rate is slightly higher due to additional meta information. When event filtering is enabled, only meta information is sent to event builder 2, which generates at least 88 bytes per event.

³The DQM receives only events from event builder 2 that have been accepted. Rejected events do not reach DQM and no ratio can be determined. Therefore the reduction by DQM corresponds to the event size reduction.



Figure 6.3: Single run with ROI filtering enabled. In the lower part the fraction is shown. The high fluctuations arise from the level-1 trigger, e.g. when a subsystem is busy and additional triggers are rejected. The right plot shows the total reduction ratio for ROI filtering only for each module individually calculated by DQM. The error bars represent the standard deviation.

High Rate Tests

Since outgoing data rates of all Selector nodes never reached more than 100 MB/s during physics runs, high throughput cannot be tested this way. Additionally the internal connection between DHE and DHC runs only at half link rate for the time being, which prevents to evaluate the performance of ONSEN at full load. High throughput tests had already been performed at the laboratory in Gießen with dummy data generated on old xFP cards. Despite this, it has never been tested with the global Belle II DAQ and actual PXD modules at KEK.

In November 2020 high rate tests were performed several times without event and ROI filtering. A $0.3\,\%$ occupancy test pattern was uploaded to the DHP memory and a $30\,\rm kHz$ Poisson trigger was used. The first 2 hours of run time are plotted in figure 6.4 demonstrating output rates above $1\,\rm GB/s$ at low memory level.

In addition, other tests have been conducted with up to $10\,\%$ occupancy pattern in order to verify the back pressure behavior. It was never observed that events had been dropped by ONSEN, where the reason could not be traced down to faulty input data.



Figure 6.4: First 2 hours of runtime of a high rate test run with 0.3% occupancy test pattern at 30 kHz. With a maximum delay of 300 ms between arrival of DHH and HLT data, the memory fill level stays at 1%. The fill level may be different with higher input rates and filtering.

Truncated Events by DHP

The previous results demonstrated ONSEN's stability in different scenarios. In addition to its main filtering purpose, data analysis features were added, which have proven their necessity multiple times. One of these features is the analysis of the common mode value within the PXD data stream. The data analysis extracts the number of events with a common mode value of 63, set by the DHP in case of full outgoing FIFOs. This strongly depends on the zero-suppressed hit count.

Reasons for high hit counts can be attributed to several factors. One common cause is poor injection, which can result in burst events flooding the sensitive area. Since increased occupancy cannot be avoided completely, a trigger veto is applied after the injection. Its length can be increased to cut out irrelevant information, but must remain small enough to guarantee a low dead time for the full detector.

Figure 6.5 shows an overview of truncated events in each physics run. Very high value counts are more likely to be caused by PXD module issues, where noisy regions cause data truncation at DHP level. An example is the first peak in June 2020, where the average occupancy in module 1.1.2 is roughly 8 times higher compared to other modules due to a problem in one of the six switcher regions. A hitmap of this specific behavior is given in figure 6.6.

Sine the beginning of monitoring $2\,472\,019$ truncated events have been detected by ONSEN and effectively supported the experts of the SuperKEKB accelerator to better understand background conditions and parameter configuration. It is a necessity to note that there is no guarantee that the common mode value in the DHP data stream is set to 63 by the DHP's internal state machine. Additionally, to



Figure 6.5: Truncated events for each run in Phase 3. In the last month of 2022 a slight upwards trend is visible, when the accelerator was pushed to its limits.



Figure 6.6: Hitmap of module 1.1.2 during a single run. The values are clipped at the 98th percentile to reveal the remaining hits. The occupancy of this module was roughly 8 times higher than other modules. Only the first 100 000 events are included. The ticks on both axes mark the geometrical borders of the ASICs. Similar behavior was also observed in other modules due to persistent radiation damage.

limit processing time the DHE may truncate events. Therefore, the total number of truncated events may be higher than anticipated.

Module Occupancy

A high impact analysis is the extraction of pixel counts. It is one of the main sources of information used by the PXD shifter to judge the PXD's behavior and performance. The pixel count is divided by the total pixels of a sensor and presented as average occupancy. The values are extracted at 2 Hz, which provides high time resolution and makes it possible to see beam changes immediately.

At higher trigger rates, maximum occupancy plays a more significant role. It allows deeper insights into beam conditions without values being too heavily smeared by averaging. Figure 6.7 plots both values. The maximum occupancy is averaged over each run. Otherwise in most cases only the truncation level of the DHE would be visible. This was increased from 8 % to 12 % in 2021 to allow smooth operation of noisy modules.



Figure 6.7: Average and maximum occupancy for each run in Phase 3. The PXD DAQ is designed to handle up to $2.5\,\%$ average occupancy.

Module Behavior

While the average occupancy of all modules is still moderate, each module behaves differently depending on its location, calibration and characteristics. In order to get a better impression of PXD modules, figure 6.8 shows the hitmaps of all inner modules side-by-side. In the upper row, the modules in forward direction are shown, in the lower area those in backward direction. The y-axis corresponds to the beam direction, while the x-axis corresponds to the phi angle, though not linearly. Module 1.3.2 (BW Sensor 3) was already problematic before installation and did not provide any useful data, so it remained switched off.

In Phase 3, the accelerator's stability and control were repeatedly tested. Small parameter adjustments and catastrophic beam abort events can have devastating consequences for the sensitive area and readout electronics of all detectors. Even the pixel detector suffered irreparable damage in some cases due to its closeness to the interaction region. However, not all performance degradation is caused by radiation.

The higher hit rates in sensors 1, 2 and 8 are due to their orientation in +X, i.e. to the outside of the accelerator ring. Sensors 4, 5 and 6 are opposite and thus receive fewer hits on average.

Also noticeable is the area at module 1.3.1 with an increased hit rate in the first switcher. It is apparent that the switcher has become corrupt and introduces noise at several gates.

Furthermore, gates are found in all modules without hits. Switchers can get damaged by beam incidents and so-called *dead gates* may occur. Such a change does not have to be permanent, and previously dead gates may provide valuable information after voltage adjustments. A similar behavior applies to drainlines, as they occur more often in module 1.4.1.

In some modules an additional ring structure can be seen. This is introduced by module manufacturing and is usually compensated for with offset and voltage



Figure 6.8: Hitmaps showing the typical hit distribution of all inner modules in a single run. Artifacts are explained in the text. Forward modules are at the top, while backward modules are at the bottom. Sensor 1 is located in the +X direction, which points to the outside of the accelerator tunnel.

calibrations. Due to the total radiation dose of the sensitive area, the operating voltages must be adjusted regularly. In this process, the compensation effect is reduced and the ring structures become more prominent.

A lower hit rate can usually be found at the right edge of each module due to the modules' geometric arrangement, which places one part of the sensor further away from the interaction point. A small portion of the sensor is also covered by the switchers of a neighboring module, making low energy photons unable to pass through.

Single Event Upsets

While the sensitive area is the most critical part, the mounted ASICs are exposed to radiation as well. The readout ASIC registers are protected against single event upsets (SEU) by triple redundancy. Therefore an SEU is unlikely but possible. The behavior was investigated by reading back the DHP registers via JTAG periodically. Since 2020 register changes have been recorded and summarized in figure 6.9.

Some SEUs could be correlated with burst events. In general, the trend follows luminosity, while extrapolation to design values estimates even higher rates. Fortunately the detection mechanism is capable of reverting the DHP register changes since 2022. However reverting changes in the DHP memory, where the pedestal, offsets and switcher sequence are stored, is still under discussion. DCD registers cannot be read out, instead blindly rewriting is the only option. During tests this rewrite procedure introduced noise in individual modules and it was decided to keep it off.



Figure 6.9: Accumulated detected SEUs in DHPs summed up over all modules.

Module Recovery

Radiation is a significant concern not only within the Belle II detector but also outside of it. Specifically, the overvoltage protection board inside the power supply situated on top of Belle II is sensitive to neutron radiation. Misbehavior in this circuit triggers emergency shutdowns. Data taking and injection is blocked until the module is powered and configured again.

Until 2019, the recovery process for single modules was manually done by the PXD shifter. This workflow was fully automated in 2020, drastically reducing PXD downtime and allowing the Belle II detector to operate more efficiently. Since 2020, 280 recoveries have been performed without shifter action, each requiring less than 3 minutes. A total of 32 of these were canceled by beam aborts causing all PXD modules to shut down. Furthermore, future changes to the slow control are foreseen to allow for continuous injection during the recovery phase. This could further enhance the Belle II detector's efficiency and effectiveness.

Missing Events

Due to the scope of ONSEN's slow control, an error in the data processing of the HLT could be identified. The decisive factor was the check of the memory level of the Selector nodes. Under normal circumstances, ONSEN receives an associated package from HLT for each event of the PXD. This ensures that all PXD events are read from memory and the memory level drops to zero when data taking stops. The memory level check was introduced to ensure that there is sufficient buffer capacity available in the following run, as there is no provision for clearing and initializing the memory at this time⁴. This resulted in a non-negligible number of unprocessed PXD events being located in memory. The cause was finally found in the underlying HLT framework (basf2) [80]. The *Event of Doom Buster* module, introduced in 2019, limits processing times when reconstructing tracks. In doing so, the entire event was silently discarded within HLT and no information was sent to ONSEN and EB2. After the module's misbehavior has been corrected, only occasional PXD events remain in ONSEN's memory, such as when an HLT worker process dies. Additional

monitoring in the ROISENDER software were implemented to detect missing HLT events already during the run.

6.2 Physics Results

Physics publications based on the Belle II data set have already been published, showcasing their significant contributions to the scientific community. One area that stands out is the measurement of particle lifetimes, which depends on reconstruction of displaced decay vertices. This would not have been possible without the remarkable performance and stability of the PXD resulting in a data set of 426 fb^{-1} . The PXD vertex resolution allows for the most precise particle lifetime measurements available consistent with recent measurements of the LHCb experiment at CERN [81, 82]. Table 6.2 lists the most recent lifetime measurements.

Particle	Measured lifetime [fs]	Reference
$\begin{array}{c} \Lambda_c^+ \\ \Omega_c^0 \end{array}$	$\begin{array}{rrrr} 203.2 \pm & 0.9 \pm & 0.8 \\ 243 & \pm 48 & \pm 11 \end{array}$	[83] [84]
$\begin{array}{c} D^0 \\ D^+ \\ D^+_s \end{array}$	$\begin{array}{rrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrrr$	[85] [85] [86]

Table 6.2: Latest charmed baryon and mesons lifetime measurements of Belle II, where the first uncertainty is statistical and the second is systematic.

Figure 6.10 displays a single event that fits the used decay channel for D^+ mesons from B meson decays in [85]. Based on the charge the trajectory is bent in the opposing direction.

⁴By design the transition from READY to RUNNING was estimated to happen within a few seconds. This is not enough time to initialize the full ONSEN memory. Therefore, it is only reset during LOADING, while HLT initializes the entire detector's geometry.



Figure 6.10: Exemplary event display from real data. Multiple reconstructed curves are shown with hits in CDC and ECL. The axis in both graphs on the right are in cm.

Chapter 7

Summary and Outlook

The work presented in this thesis has been focused around the development and integration of the ONSEN data acquisition system of the pixel detector (PXD) at the Belle II experiment and its commissioning and operation during physics data taking from 2019 to 2022. Further, PXD's monitoring system was designed, implemented and operated. During this period the accelerator reached the highest luminosity to date of 4.71×10^{34} cm⁻² s⁻¹ in June 2022, at level-1 trigger rates of more than 8 kHz. This superseded the luminosity record of the LHC by about a factor of 2 [87]. In total the Belle II detector recorded 426 fb⁻¹, which corresponds to about 350 million B meson events.

Firstly, the work on ONSEN's firmware and software is described. Improvements in data handling and data analysis were made to provide real-time information during data taking and give feedback about data quality to the shifter.

- With a time resolution of 0.5 seconds, module occupancy extracted on the fly by ONSEN has proven to be the most accurate and reliable source.
- An average hit occupancy of $0.25\,\%$ was observed at the end of the run period.
- Single high background events with more than 10 % occupancy pushed the detector's readout ASICs over their limits. Data truncation has been detected by ONSEN in over 2.4 million events supporting beam operation and injection parameter choice successfully.
- The monitoring of the ONSEN system indicated that nearly 300 TB of zero-suppressed raw data was produced by PXD modules and analyzed by ONSEN during physics data taking. Data was successfully sent to EB2 for storage after ONSEN filtering, which removed 60 % of the overall data.
- In test runs at KEK with the full Belle II DAQ chain, ONSEN's stable operation was demonstrated with trigger rates of over 35 kHz and total incoming data rates of up to 4.8 GB/s. This corresponds to the maximum possible data rate of 4 (out of 8) DHCs running with half the bandwidth between DHC and DHE.

- The load balancing mechanism between DHC and ONSEN Selector nodes operated without incident for the entire period.
- ROI filtering was validated for several runs with a reduction factor of up to 10 and is prepared for higher luminosity operation beyond 2024.

In light of the installation of the fully equipped PXD during the long shutdown, ONSEN is well prepared to resume operation.

Secondly, the slow control system of the entire PXD underwent improvements to better fit its operational needs. Permanently changing beam conditions pushed PXD module operation into unknown territory.

- Unexpected downtimes of single PXD modules burdened shifter supervision until recovery was fully automated, vastly improving data taking efficiency about 250 times.
- Monitoring of the ASIC's internal registers including correction of radiationinduced register flips allowed continuous data taking. More than 35 single event upsets were detected in a period of 2.5 years and in the last run period all were automatically recovered.
- Increased radiation exposure resulted in unexpected behavior in certain electrical channels of PXD modules. The ramping sequence received multiple improvements to counteract these while ensuring steady currents within safe margins. As a side effect, this sped up the overall ramping process to allow for faster recovery.
- Pedestal calibration has been automated to the point where no shifter action is required and data taking is almost not impacted. Most calibration methods have been unified, handling the entire PXD ecosystem and reducing interference from various monitoring mechanisms.

All modifications to the software were made while keeping PXD modules' safety top priority. With growing complexity, stability and speed have become another significant aspect to ensure the highest possible uptime during valuable beam operation. All improvements are scalable and will provide the same level of reliability for the full PXD system with 40 modules, scheduled to be operational by the end of 2023.

From a usability point of view, the PXD is in excellent condition and future adjustments will be minimal. Since the fresh PXD modules will experience a higher dose during their lifetime, their behavior cannot be completely forecasted in advance, and further adjustments may be inevidable.

Acknowledgements

Let me start by expressing my sincere gratitude to my supervisor Sören Lange for all the support and professional and personal assistance that he provided throughout my PhD.

Secondly, I would like to thank Wolfgang Kühn and Claudia Höhne for their continuous support and giving me to the opportunity to work on such an interesting and motivating research topic.

Also, I would like to thank Thomas Geßler for his constructive feedback, which was extremely helpful. He provides invaluable guidance and expertise regarding firmware development.

Also, I would like to thank the entire working group for their assistance in addressing topics concerning both on-topic and off-topic questions, and in particular, Klemens Lautenbach and Leonard Koch for numerous barbecues and other fun activities that happened along the way during the course of our work together. Special thanks to Dennis Getzkow, with whom I have worked together on several projects in the same office and at KEK.

Thanks to everyone at the University of Gießen, KEK, and the Belle II experiment for providing me with the resources and support to do my research.

Finally, I wish to thank all the friends I made before and during my PhD for making these past years unforgettable. I would like to thank my parents and my brother for all their support and encouragement throughout my studies.

References

- Chen-Ning Yang and Robert L. Mills. "Conservation of Isotopic Spin and Isotopic Gauge Invariance." *Phys. Rev.* 96 (1954). Ed. by Jong-Ping Hsu and D. Fine, 191–195.
 DOI: 10.1103/PhysRev.96.191.
- S. L. Glashow. "Partial Symmetries of Weak Interactions." *Nucl. Phys.* 22 (1961), 579–588.
 DOI: 10.1016/0029-5582(61)90469-2.
- [3] Steven Weinberg. "A Model of Leptons." *Phys. Rev. Lett.* 19 (1967), 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.
- [4] F. J. Hasert et al. "Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment." *Phys. Lett. B* 46 (1973), 138–140.
 DOI: 10.1016/0370-2693(73)90499-1.
- [5] W. Pauli. "Relativistic Field Theories of Elementary Particles." *Rev. Mod. Phys.* 13 (1941), 203–232.
 DOI: 10.1103/RevModPhys.13.203.
- S. L. Glashow, J. Iliopoulos, and L. Maiani. "Weak Interactions with Lepton-Hadron Symmetry." *Phys. Rev. D* 2 (1970), 1285–1292.
 DOI: 10.1103/PhysRevD.2.1285.
- [7] M. Tanabashi et al. "Review of Particle Physics." *Phys. Rev. D* 98.3 (2018), 030001.
 DOI: 10.1103/PhysRevD.98.030001.
- [8] Mark Thomson. Modern particle physics. New York: Cambridge University Press, 2013. ISBN: 978-1-107-03426-6.
 DOI: 10.1017/CBO9781139525367.
- [9] S. K. Choi et al. "Observation of a narrow charmonium-like state in exclusive $B^{\pm} \rightarrow K^{\pm}\pi^{+}\pi^{-}J/\psi$ decays." *Phys. Rev. Lett.* 91 (2003), 262001. DOI: 10.1103/PhysRevLett.91.262001. arXiv: hep-ex/0309032.

- [10] Bernard Aubert et al. "Study of the $B \rightarrow J/\psi K^- \pi^+ \pi^-$ decay and measurement of the $B \rightarrow X(3872)K^-$ branching fraction." *Phys. Rev. D* 71 (2005), 071103. DOI: 10.1103/PhysRevD.71.071103. arXiv: hep-ex/0406022.
- [11] M. Ablikim et al. "Observation of $e^+e^- \rightarrow \gamma X(3872)$ at BESIII." *Phys. Rev. Lett.* 112.9 (2014), 092001. DOI: 10.1103/PhysRevLett.112.092001. arXiv: 1310.4101 [hep-ex].
- [12] Makoto Kobayashi and Toshihide Maskawa. "CP Violation in the Renormalizable Theory of Weak Interaction." *Prog. Theor. Phys.* 49 (1973), 652–657.
 DOI: 10.1143/PTP.49.652.
- [13] A. J. Bevan et al. "The Physics of the B Factories." *Eur. Phys. J. C* 74 (2014), 3026.
 DOI: 10.1140/epjc/s10052-014-3026-9.
 arXiv: 1406.6311 [hep-ex].
- [14] Albert Einstein, Boris Podolsky, and Nathan Rosen. "Can quantum mechanical description of physical reality be considered complete?" *Phys. Rev.* 47 (1935), 777–780.
 DOI: 10.1103/PhysRev.47.777.
- [15] Ikaros I. Y. Bigi and A. I. Sanda. "Notes on the Observability of CP Violations in B Decays." *Nucl. Phys. B* 193 (1981), 85–108.
 DOI: 10.1016/0550-3213(81)90519-8.
- [16] Kazuo Abe et al. "Observation of large CP violation in the neutral *B* meson system." *Phys. Rev. Lett.* 87 (2001), 091802.
 DOI: 10.1103/PhysRevLett.87.091802.
 arXiv: hep-ex/0107061.
- Bernard Aubert et al. "Observation of CP violation in the B⁰ meson system." *Phys. Rev. Lett.* 87 (2001), 091801.
 DOI: 10.1103/PhysRevLett.87.091801. arXiv: hep-ex/0107013.
- [18] Lincoln Wolfenstein. "Parametrization of the Kobayashi-Maskawa Matrix." *Phys. Rev. Lett.* 51 (1983), 1945.
 DOI: 10.1103/PhysRevLett.51.1945.
- [19] I. Adachi et al. "Precise measurement of the CP violation parameter sin $2\phi_1$ in $B^0 \rightarrow (c\bar{c})K^0$ decays." *Phys. Rev. Lett.* 108 (2012), 171802. DOI: 10.1103/PhysRevLett.108.171802. arXiv: 1201.4643 [hep-ex].
- J. Charles et al. "CP violation and the CKM matrix: Assessing the impact of the asymmetric *B* factories." *Eur. Phys. J. C* 41.1 (2005), 1–131.
 DOI: 10.1140/epjc/s2005-02169-1. arXiv: hep-ph/0406184.

- [21] T. Abe et al. "Belle II Technical Design Report" (Nov. 2010).
 DOI: 10.48550/arXiv.1011.0352.
 arXiv: 1011.0352 [physics.ins-det].
- [22] CERN Courier. KEKB Breaks Luminosity Record. CERN Courier. June 8, 2009. URL: https://cerncourier.com/a/kekb-breaksluminosity-record/(visited on 05/16/2023).
- [23] Demin Zhou et al. "Luminosity performance of SuperKEKB" (June 2023). arXiv: 2306.02692 [physics.acc-ph].
- [24] M. Yamaga et al. "RPC systems for BELLE detector at KEKB." *Nucl. Instrum. Meth. A* 456 (2000). Ed. by G. Grancagnolo et al., 109–112.
 DOI: 10.1016/S0168-9002(00)00973-6.
- [25] T. Aushev et al. "A scintillator based endcap K_L and muon detector for the Belle II experiment." *Nucl. Instrum. Meth. A* 789 (2015), 134–142.
 DOI: 10.1016/j.nima.2015.03.060.
 arXiv: 1406.3267 [physics.ins-det].
- [26] A. Kuzmin. "Endcap calorimeter for SuperBelle based on pure CsI crystals." Nucl. Instrum. Meth. A 623 (2010). Ed. by Hiroyuki Iwasaki, Takeshi K. Komatsubara, and Yasuhiro Sugimoto, 252–254.
 DOI: 10.1016/j.nima.2010.02.212.
- [27] R. Pestotnik et al. "Aerogel RICH for forward PID at Belle II." *Nucl. Instrum. Meth. A* 732 (2013). Ed. by T. Bergauer et al., 371–374.
 DOI: 10.1016/j.nima.2013.06.080.
- [28] Kenji Inami. "TOP counter for particle identification at the Belle II experiment." *Nucl. Instrum. Meth. A* 766 (2014). Ed. by T. Sumiyoshi et al., 5–8. DOI: 10.1016/j.nima.2014.07.006.
- [29] Nanae Taniguchi et al. "All-in-one readout electronics for the Belle-II Central Drift Chamber." *Nucl. Instrum. Meth. A* 732 (2013). Ed. by T. Bergauer et al., 540–542.
 - DOI: 10.1016/j.nima.2013.06.096.
- [30] M. Friedl et al. "The Belle II Silicon Vertex Detector." *Phys. Procedia* 37 (2012).
 Ed. by Ted Liu, 867–873.
 DOI: 10.1016/j.phpro.2012.02.428.
- [31] M. J. French et al. "Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker." *Nucl. Instrum. Meth. A* 466 (2001). Ed. by T. Ohsugi et al., 359–365.
 DOI: 10.1016/S0168-9002(01)00589-7.
- [32] Andreas Moll. "Comprehensive study of the background for the Pixel Vertex Detector at Belle II." PhD thesis. Munich U., 2015.
 DOI: 10.5282/edoc.19106.
- [33] A. Natochii et al. "Beam background expectations for Belle II at SuperKEKB." Snowmass 2021. Mar. 2022. arXiv: 2203.05731 [hep-ex].

- [34] W. Neeser et al. "The DEPFET pixel BIOSCOPE." *IEEE Trans. Nucl. Sci.* 47 (2000), 1246–1250.
 DOI: 10.1109/23.856581.
- [35] Hans Kruger. "Front-end electronics for DEPFET pixel detectors at SuperBelle (BELLE II)." Nucl. Instrum. Meth. A 617 (2010). Ed. by Giorgio Chiarelli et al., 337–341.
 DOI: 10.1016/j.nima.2009.10.042.
- [36] J. Knopf et al. "A 256 channel 8-Bit current digitizer ASIC for the Belle-II PXD." JINST 6 (2011), C01085.
 DOI: 10.1088/1748-0221/6/01/C01085.
- [37] M. Lemarenko et al. "Test results of the data handling processor for the DEPFET pixel vertex detector." *JINST* 8 (2013), C01032.
 DOI: 10.1088/1748-0221/8/01/C01032.
- [38] Yoshihito Iwasaki et al. "Level 1 trigger system for the Belle II experiment." *IEEE Trans. Nucl. Sci.* 58 (2011). Ed. by Sascha Marc Schmeling, 1807–1815. DOI: 10.1109/TNS.2011.2119329.
- [39] M. Nakao. "Timing distribution for the Belle II data acquistion system." *JINST* 7 (2012), C01028.
 DOI: 10.1088/1748-0221/7/01/C01028.
- [40] M. Nakao et al. "Data acquisition system for Belle II." JINST 5 (2010), C12004.
 DOI: 10.1088/1748-0221/5/12/C12004.
- [41] T. Higuchi et al. "Development of a PCI based data acquisition platform for high intensity accelerator experiments." *eConf* C0303241 (2003), TUGT004. arXiv: hep-ex/0305088.
- [42] R. Itoh et al. "Data flow and high level trigger of Belle II DAQ system." *IEEE Trans. Nucl. Sci.* 60 (2013). Ed. by Dora Merelli, 3720–3724.
 DOI: 10.1109/RTC.2012.6418174.
- [43] Matthew Barrett et al. "The Belle II Online–Offline Data Operations System." *Comput. Softw. Big Sci.* 5.1 (2021), 1. DOI: 10.1007/s41781-020-00045-9.
- [44] R. Brun and F. Rademakers. "ROOT: An object oriented data analysis framework." *Nucl. Instrum. Meth. A* 389 (1997). Ed. by M. Werlen and D. Perret-Gallix, 81–86.
 DOI: 10.1016/S0168-9002(97)00048-X.

[45] Dmytro Levit et al. "FPGA based data read-out system of the Belle II pixel detector" (2014), 7097505.
 DOI: 10.1109/RTC.2014.7097505.
 arXiv: 1406.3864 [physics.ins-det].

[46] Stefan Huber et al. "Performance of the Data-Handling Hub Readout System for the Belle II Pixel Detector." *IEEE Trans. Nucl. Sci.* 68.8 (2021), 1961–1967. DOI: 10.1109/TNS.2021.3083720. arXiv: 2011.00067 [physics.ins-det].
- [47] Michael Schnell. "Development of an FPGA-based Data Reduction System for the Belle II DEPFET Pixel Detector." PhD thesis. Bonn U., 2015.
- [48] Bruno Deschamps. "Development and Deployment of DATCON, an FPGAbased Data Reduction System for the Belle II Pixel Detector." PhD thesis. U. Bonn (main), 2022.
- [49] Björn Spruck et al. "The Belle II Pixel Detector Data Acquisition and Reduction System." *IEEE Trans. Nucl. Sci.* 60.5 (2013), 3709–3713.
 DOI: 10.1109/TNS.2013.2281571.
- [50] David Münchow et al. "The data acquisition system of the Belle II Pixel Detector." *JINST* 9 (2014), C08009.
 DOI: 10.1088/1748-0221/9/08/C08009.
- [51] Thomas Geßler et al. "The ONSEN Data Reduction System for the Belle II Pixel Detector." *IEEE Trans. Nucl. Sci.* 62.3 (2015), 1149–1154.
 DOI: 10.1109/TNS.2015.2414713. arXiv: 1406.4028 [physics.ins-det].
- [52] PICMG. PICMG 3.0 R3.0, AdvancedTCA Base Specification. 2008.
- [53] PICMG. AMC.0 R2.0, AdvancedMC Mezzanine Module Specification. Version 2.0. 2006.
- [54] PICMG. MTCA.0 R1.0, MicroTCA Specification. 2006.
- [55] Xilinx, Inc. DS112. Virtex-4 Family Overview. Version 3.1. 2010. URL: https://docs.xilinx.com/v/u/en-US/plbv46_ slave_single.
- [56] Xilinx, Inc. DS112. Virtex-5 Family Overview. Version 5.1. 2015. URL: https://docs.xilinx.com/v/u/en-US/ds100.
- [57] Xilinx, Inc. SP006. LocalLink Interface Specification. Version 2.0. 2005. URL: http://www.xilinx.com/aurora/aurora_member/ sp006.pdf.
- [58] Xilinx, Inc. UG683. EDK Concepts, Tools, and Techniques. Version 14.7. 2013. URL: http://www.xilinx.com/support/documentation/ sw_manuals/xilinx14_7/edk_ctt.pdf.
- [59] Xilinx, Inc. DS561. PLBv46 Slave Single, Data Sheet. Version 1.3. 2010. URL: https://docs.xilinx.com/v/u/en-US/ds112.
- [60] Buildroot Asociation. *Buildroot*. URL: https://buildroot.org/.
- [61] TIS Committe. Tool Interface Standard (TIS) Portable Formats Specification. Oct. 1993.
- [62] Thomas Gessler. "Development of FPGA-Based Algorithms for the Data Acquisition of the Belle II Pixel Detector." PhD thesis. Giessen U., 2015.

- [63] J.O. Hill. "Channel access: A software bus for the LAACS." Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 293.1 (1990), 352–355. ISSN: 0168-9002.
 DOI: 10.1016/0168-9002(90)91459-0.
 URL: https://www.sciencedirect.com/science/article/ pii/0168900290914590.
- [64] L R Dalesio, A J Kozubal, and M R Kraimer. "EPICS architecture" (Jan. 1991). URL: https://www.osti.gov/biblio/6110347.
- [65] EPICS. Control System Studio. URL: https://controlsystemstudio.org/(visited on 05/18/2022).
- [66] Christopher Slominski. "A MySQL Based EPICS Archiver" (Oct. 2009). URL: https://www.osti.gov/biblio/1023592.
- [67] Xilinx, Inc. UG633. PlanAhead Methodology Guide. Version 14.6. 2008. URL: https://docs.xilinx.com/v/u/en-US/PlanAhead10-1_MethodologyGuide.
- [68] Michael Ritzert. "Interfacing EPICS to the Widespread Platform Management Interface IPMI." 15th International Conference on Accelerator and Large Experimental Physics Control Systems. 2015, WEPGF024.
 DOI: 10.18429/JACOW-ICALEPCS2015-WEPGF024.
- [69] Xilinx, Inc. SPOO. Aurora 8B/10B Protocol Specification. Version 2.3. 2014. URL: https://docs.xilinx.com/v/u/en-US/aurora_ 8b10b_protocol_spec_sp002.
- [70] Matthew Newville. PyEpics Epics Channel Access for Python. URL: https://cars9.uchicago.edu/software/python/ pyepics3/index.html (visited on 12/16/2022).
- [71] Wang Xiaoqiang. PCASpy Documentation. URL: https://pcaspy.readthedocs.io/en/latest/ (visited on 12/16/2022).
- [72] Holger Krekel and pytest-dev team. *Pytest Documentation*. pytest: helps you write better programs.
 URL: https://docs.pytest.org/en/7.0.x/ (visited on 12/19/2022).
- [73] Atlassian. Collaboration Software for Software, IT and Business Teams. Atlassian. URL: https://www.atlassian.com(visited on 12/19/2022).
- [74] GitLab B.V. GitLab The One DevOps Platform. URL: https://about.gitlab.com/ (visited on 12/19/2022).
- [75] Harrison Schreeck. "Commissioning and first data taking experience with the Belle II pixel vertex detector." PhD thesis. Gottingen U., II. Phys. Inst., 2020. DOI: 10.53846/goediss-8060.

- [76] Docker Inc. Docker: Accelerated, Containerized Application Development. May 10, 2022.
 URL: https://www.docker.com/ (visited on 12/20/2022).
- [77] Stefan Behnel et al. "Cython: The Best of Both Worlds." *Computing in Science & Engineering* 13.2 (2 2011), 31–39. ISSN: 1558-366X.
 DOI: 10.1109/MCSE.2010.118.
- [78] Charles R. Harris et al. "Array programming with NumPy." *Nature* 585.7825 (2020), 357–362.
 DOI: 10.1038/s41586-020-2649-2.
 arXiv: 2006.10256 [cs.MS].
- [79] John D. Hunter. "Matplotlib: A 2D Graphics Environment." *Comput. Sci. Eng.* 9.3 (2007), 90–95.
 DOI: 10.1109/MCSE.2007.55.
- [80] T. Kuhr et al. "The Belle II Core Software." *Comput. Softw. Big Sci.* 3.1 (2019), 1. DOI: 10.1007/s41781-018-0017-9. arXiv: 1809.04299 [physics.comp-ph].
- [81] Roel Aaij et al. "Measurement of the lifetimes of promptly produced Ω_c^0 and Ξ_c^0 baryons." *Sci. Bull.* 67.5 (2022), 479–487. DOI: 10.1016/j.scib.2021.11.022. arXiv: 2109.01334 [hep-ex].
- [82] Roel Aaij et al. "Precision measurement of the Λ_c^+ , Ξ_c^+ and Ξ_c^0 baryon lifetimes." *Phys. Rev. D* 100.3 (2019), 032001. DOI: 10.1103/PhysRevD.100.032001. arXiv: 1906.08350 [hep-ex].
- [83] F. Abudinén et al. "Measurement of the Λ_c^+ Lifetime." *Phys. Rev. Lett.* 130.7 (2023), 071802. DOI: 10.1103/PhysRevLett.130.071802. arXiv: 2206.15227 [hep-ex].
- [84] Fernando Jesus Abudinen et al. "Measurement of the Ω_c^0 lifetime at Belle II." *Phys. Rev. D* 107.3 (2023), L031103. DOI: 10.1103/PhysRevD.107.L031103. arXiv: 2208.08573 [hep-ex].
- [85] F. Abudinén et al. "Precise measurement of the D⁰ and D⁺ lifetimes at Belle II." *Phys. Rev. Lett.* 127.21 (2021), 211801.
 DOI: 10.1103/PhysRevLett.127.211801.
 arXiv: 2108.03216 [hep-ex].
- [86] I. Adachi et al. "Precise measurement of the D⁺_s lifetime at Belle II" (June 2023).
 arXiv: 2306.00365 [hep-ex].
- [87] LHC announcement. Record luminosity: well done LHC. LHC. Nov. 13, 2017. URL: https://home.cern/news/news/accelerators/ record-luminosity-well-done-lhc (visited on 08/16/2023).